Wax

22234

ARPA Network Information Center
Stanford Research Institute
Melno Park, California 94025
NIC # 22234
Packet Radio Temporary Note #95

Network Analysis Corp.
March, 1974

PACKET RADIO SIMULATION RESULTS - I

1. INTRODUCTION

This paper describes experimental results for the Packet Radio System [1], obtained by simulation. The simulation [2], was developed as a design tool to assist in solving intractable analytic problems. Consequently, we expect to further experiment with the program to answer practical questions pertinent to system design. Two series of experiments are discussed. In the first series, some of the bottle necks in system performance are idetified. In the second series, specific questions related to system design, such as the trade offs between device range and device interference and the trade offs between single and dual data rate systems are examined. We also estimate the maximum throughput, delays, and blocking in the systems simulated. The aspects of system operation which differ from the formal presentations [1], [3], are described in the Appendix.

2. PERFORMANCE MEASURES

The packet radio system is a loss and delay system. Thus, throughput and delay measures in the usual sense are not sufficient to characterize system performance. In particular, one cannot obtain a case in which the delay is very high since above a certain delay blocking will increase faster than the delay. The blocking phenomena is desirable since it can be used as a control to limit the traffic level. The following measurements taken in the simulation are used to evaluate system performance.

Throughput: Throughput is measured per terminal slots. It is measured at the station and evaluated by counting the number of IPs* from terminals received at the station plus the number of IPs transmitted to terminals and which have been ETE acknowledged, divided by the number of terminal simulation slots. The actual throughput of the system may be slightly higher than the one measured since some of the IPs may have arrived to terminals and not yet been acknowledged, while some ETE Acks can be received at the station after the station discarded packets, these are not counted.

Delay: The following delays are measured: terminal search delay, the delay of the arrival of the first IP from the terminal to the station, the delay until the terminal receives the ETE Acks from the station, and total time spent in the

* The definition of packet types simulated is given in the Appendix.

system (terminal interaction delay). The search delay is not significant, it will usually be below one terminal slot; otherwise the terminal will most often be blocked. The delay of the arrival of the first IP to the station and the ETE Ack back to the terminal seem to be significant since they are independent of the delays in the PTP network and of the number of packets that the terminal has to receive in response. Furthermore, the ETE Ack delay is an indication that the terminal established communication with the PTP network.

Prob. Station Busy: The station is sampled during the simulation and is assumed busy if it actively receiving or transmitting. Otherwise, it is considered idle. This measure is an indication of the channel traffic at the station.

Blocking: If a terminal does not identify a specific receiver after transmitting a SP, the number of times specified, it departs from the system. We say that such a terminal is blocked. The packets considered blocked in this case are the packet from the terminal to the station and the response packets that the terminal was supposed to receive from the station. We use this count to evaluate the percentage of IPs blocked.

Digitized by the Internet Archive
in 2025 with funding from
Amateur Radio Digital Communications, Grant 151

https://archive.org/details/temporary-note-096

Total Loss: In addition to the terminals blocked, some terminals may depart from the system, after establishing communication with the station, without receiving any or all of the response packets from the station.

Terminals Remaining: This indicates the number of terminals which are still in the system when the simulation ends.

Rate of Station Response: A parameter in the simulation indicates the average number of response packets from station to terminals. When the value measured is much below the average, it indicates that the system is saturated since the offered traffic rate to the system is much higher than the maximum throughput of the system. This will also be reflected in the difference between the total loss and the blocking, as well as in the number of terminals remaining in the system.

Remarks: (1) Some of the measures give "overlapping" information, and are taken in order to obtain more detailed observations of system performance. For example, the total loss indicates the probability of blocking in the system; however, by taking the blocking measurement, one is able to determine whether the loss is due to the inefficiency in routing or because of difficulty in entering the system.

(2) The differences between offered rate, throughput and loss is in the number of terminals remaining.

3. PRELIMINARY OBSERVATIONS

The first system simulated was a Common Channel Single Data Rate system, in which the station is routing traffic as a repeater (Naive Station). We denote the system as CCSDR (NS). The system defined has a single data signalling rate for communication between terminal and repeater (or station) and in the repeater-station network; the channel is used in a half duplex mode. When the station is routing traffic as a repeater, it cannot receive packets not specifically addressed to i

In all experiments reported here, the labels of repeaters and station were preassigned. The hierarchical (diredted) labeling scheme of the system in this experiment are shown in Figure 1. Figure 2 shows the connectivity of the repeaters and station. That is, when a device transmits, all the devices connected to it by line are within an effective range and "hear" the transmission.

The objective of the first series of experiments was to observe the detailed operation of devices and the efficiency of the system. The following observations were made:

(1) The "critical hop" in the system is that between the first level repeaters and the station. This was concluded by observing the frequency at which repeaters begin to search and at which they discarded packets, and from the observation that there is no significant difference in the delay when the number of hops from the station that a packet travels is increased.

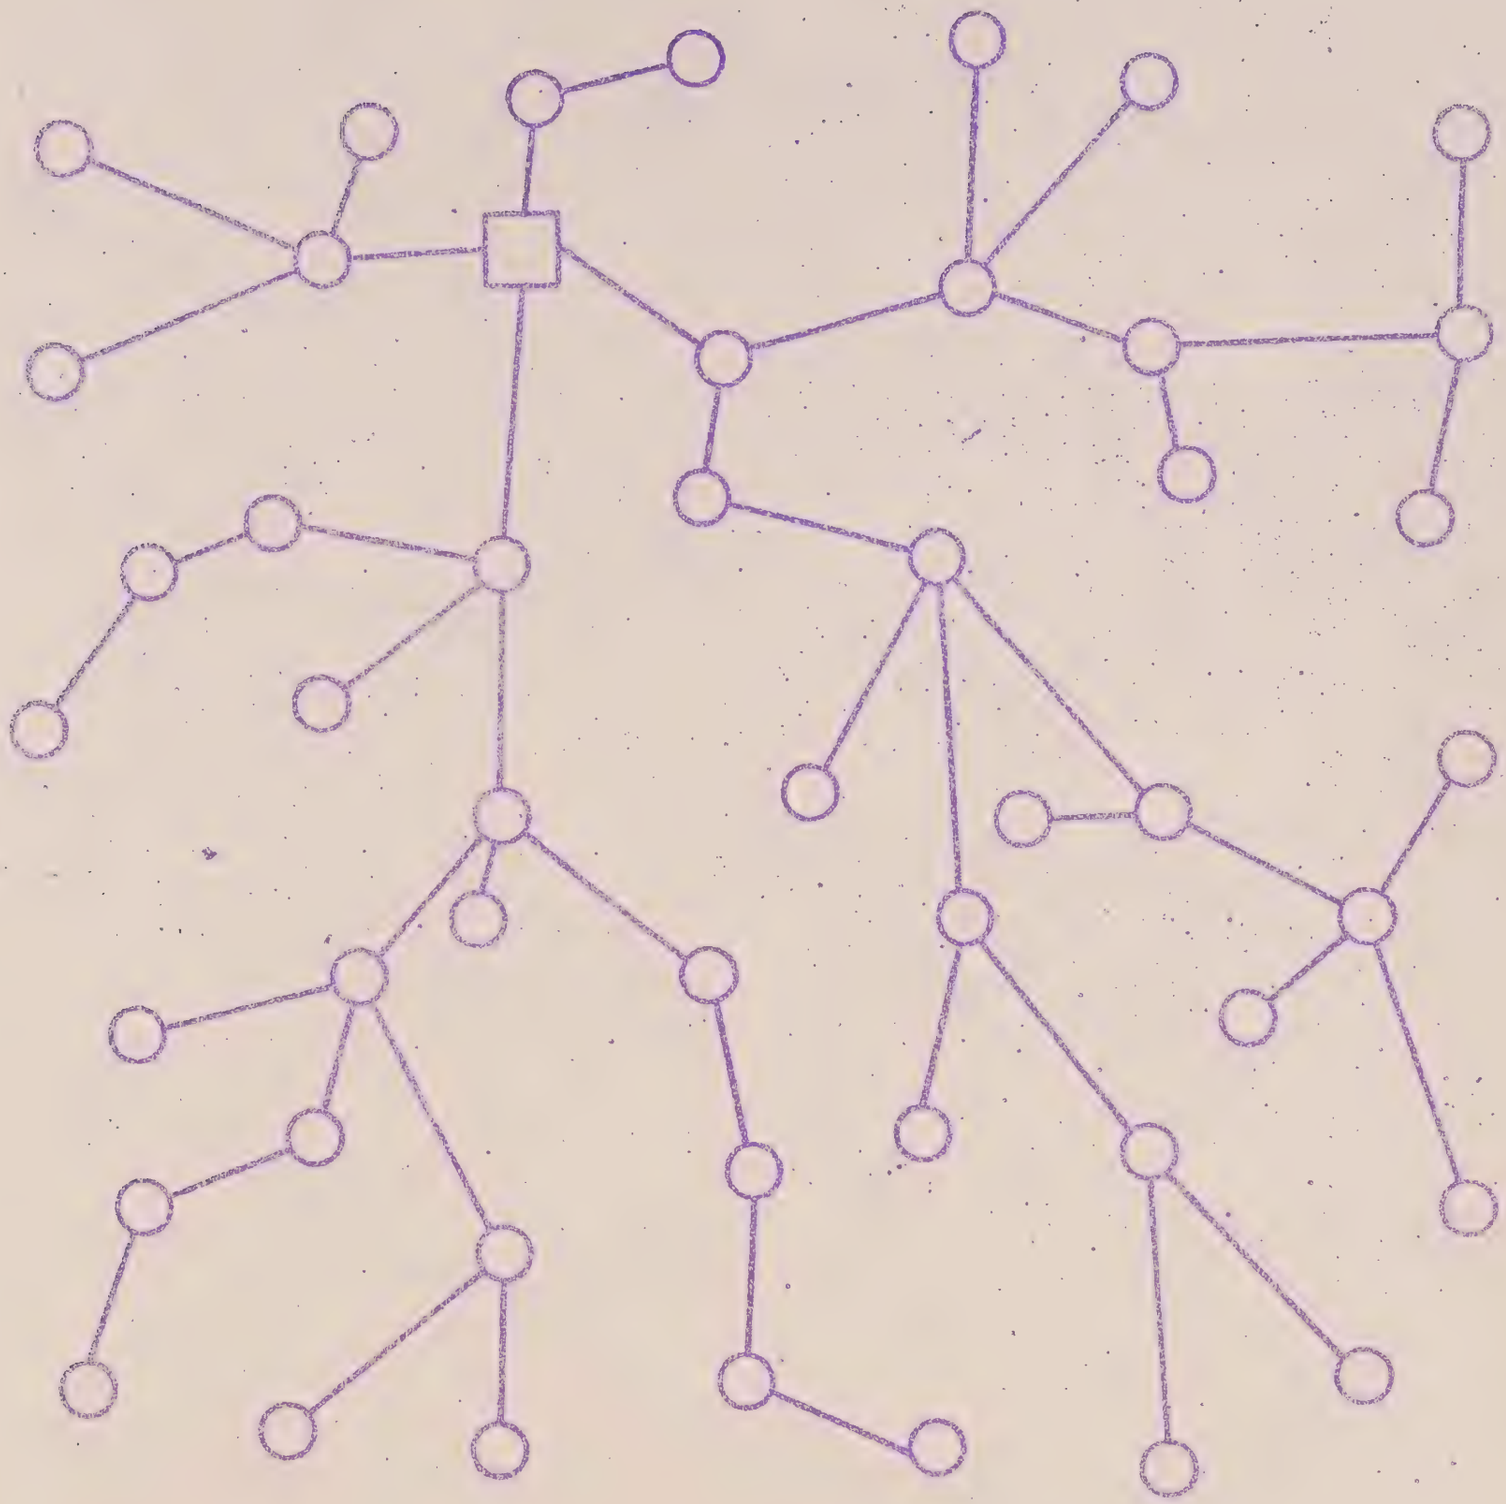

FIGURE 1

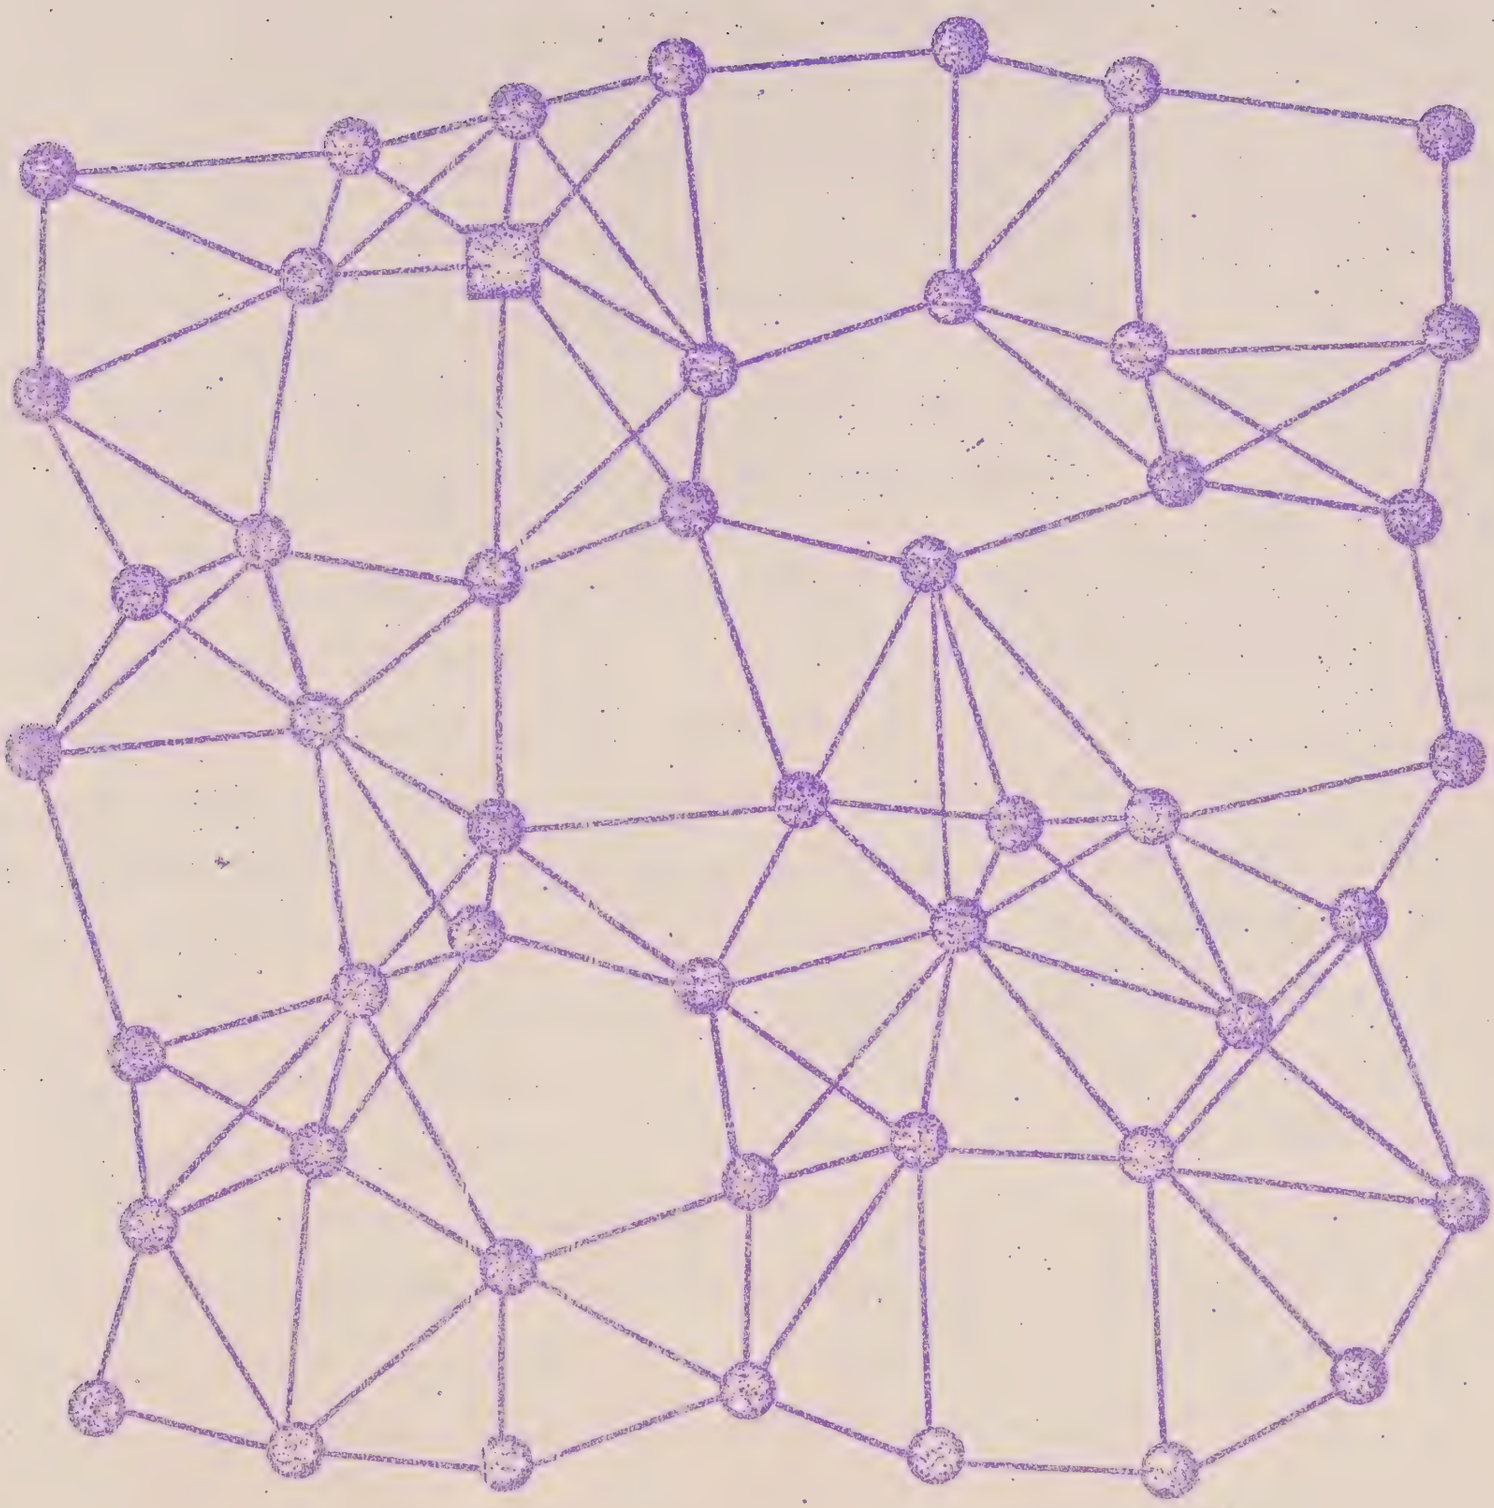

FIGURE 2

(2) There is a higher probability of end-to-end successful completion when routing from the station to a terminal than when routing from a terminal to the station. Practically, there is almost "no" difference in time delay between the delay of an information packet from the terminal arriving at the station and the time that the terminal receives an ETE Ack from the station.

(3) Many packets associated with terminals that have departed from the system are routed in the network.

The effect of improving the routing capabilities of the station can be readily observed. In particular, one can see in Figures 2 and 1 that while the connectivity of the station is 7, there are only 4 repeaters labelled from the station. Consequently, the station is busy part of the time with non-useful traffic. This situation can be improved by changing the routing of the station so that: (1) it receives any packet that it can hear and which is (eventually) addressed to it; and (2) it transmits response packets to the repeater nearest to the terminal along the routing path that it can reach. This change was implemented for all system studies subsequent to the initial experiments. Apart from the change implemented, the observation suggests that particular attention should be given to the design of the repeater network in the neighborhood of the stations. It is also noted that these repeaters have a higher power duty cycle since they handle all packets collected from other

parts of the network. The routing change made at the station enables the allocation of many more repeaters in the neighborhood of the station, than are functionally needed, without resulting an increase in the artificial traffic generated. The exact labelling of these repeaters is also not critical. One of the reasons leading to observation 2 is that the station has a higher probability, than the first level repeaters, of successful transmission over the critical hop because it is the largest user and does not interfere with its own transmissions. Theoretically, one may expect a similar conclussion when considering transmission in a section of the network in which two repeaters, one of which "homes" on the other, compete. This,however, may not be realized in the system simulated because of the limited storage available in repeaters.

Observations 2 and 3 suggest a change in the Terminal-Station protocol. The basic question is whether a terminal should release itself from the system or whether it should be released by the station. The former was initially simulated. It was observed that in many cases a terminal departed from the system after receiving an Echo to the ETE Ack for the last IP without this ETE Ack arriving at the station. This resulted in the reactivation of IP's by the station for this terminal, the routing of these packets in the net, and then the maximum number of transmissions and search by the repeater nearest to the terminal. The protocol simulated in the systems discussed later is such that the last transmission must always be from the station to the terminal. This transmission may be considered as a terminal release packet. Another change in protocol implemented is that whenever possible, the terminal

acknowledges a sequence of packets rather than individual ones, to reduce the overhead in the direction towards the station.

4. RANGE VERSUS INTERFERENCE

For the experiments discussed in the previous section, it was assumed that Repeater-Repeater range is the same as Terminal-Repeater range. This, however, is not always a realistic assumption since repeaters can be placed on elevated areas and can have more power then terminals, (especially hand held terminals). Thus, if repeaters are allocated for area coverage of terminals, the repeater range will be higher than terminal range and higher network connectivity or device interference will result.

The problem which then arrises is to determine the impact of this interference on system performance. Alternatively, one may seek to reduce repeater transmission power when transmitting in the repeater-station network. To study this issue, two CCSDR systems were simulated, one with high interference CCSDR (HI), and the other with Low Interference CCSDR (LI). The routing labels of the two systems were the same and are shown in Figure 1. The interference of the CCSDR (LI) system is shown in Figure 2 and the interference of the CCSDR (HI) system in Figure 3. Figure 3 shows only the connecivity of two devices in the network.

The results are shown in Figure 4 and Table 1. Figure 4 shows the throughput of the two systems as a function of time while Table 1 summarizes all other measures of performance. The third row of Table 1 summarizes performance of the high interference system under an improved set of repeater

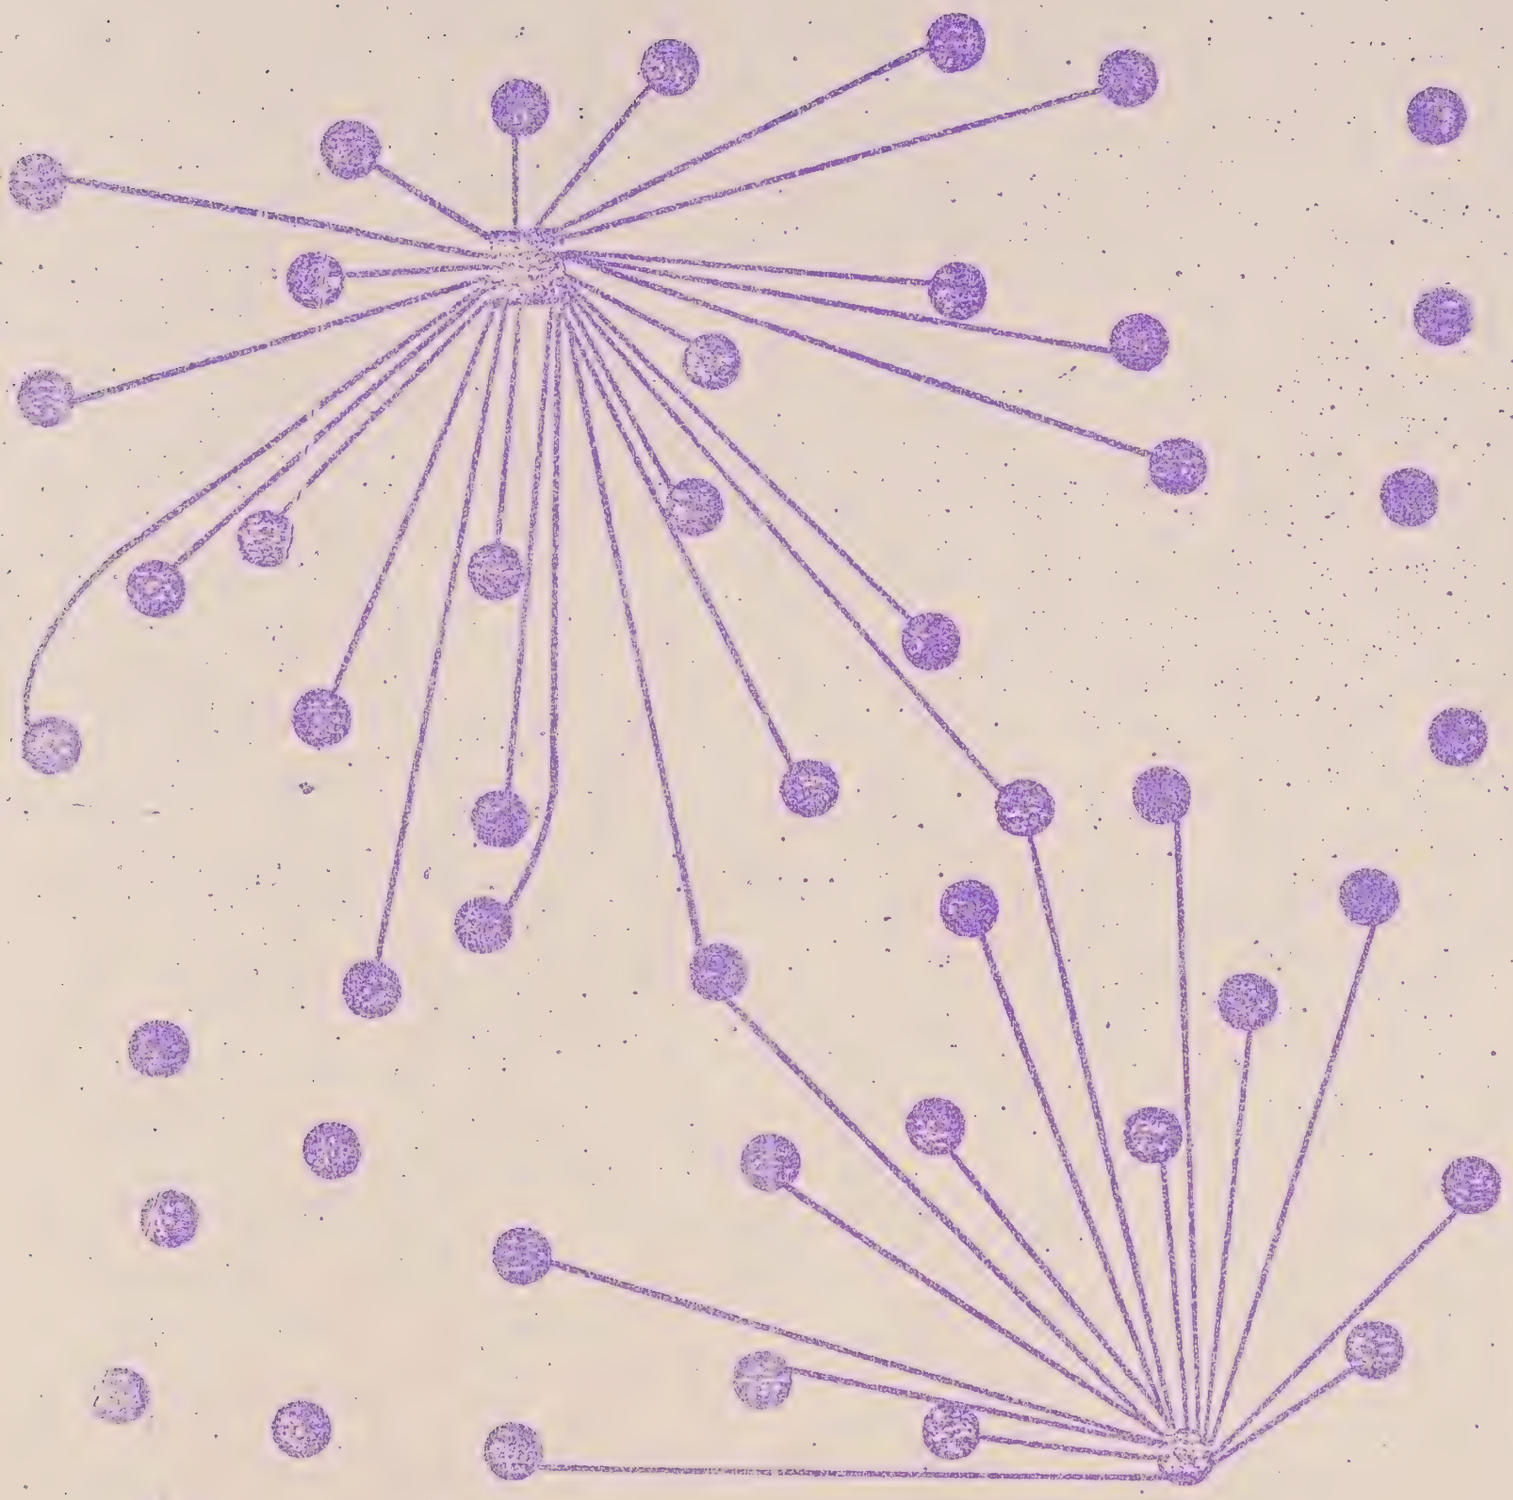

FIGURE 3

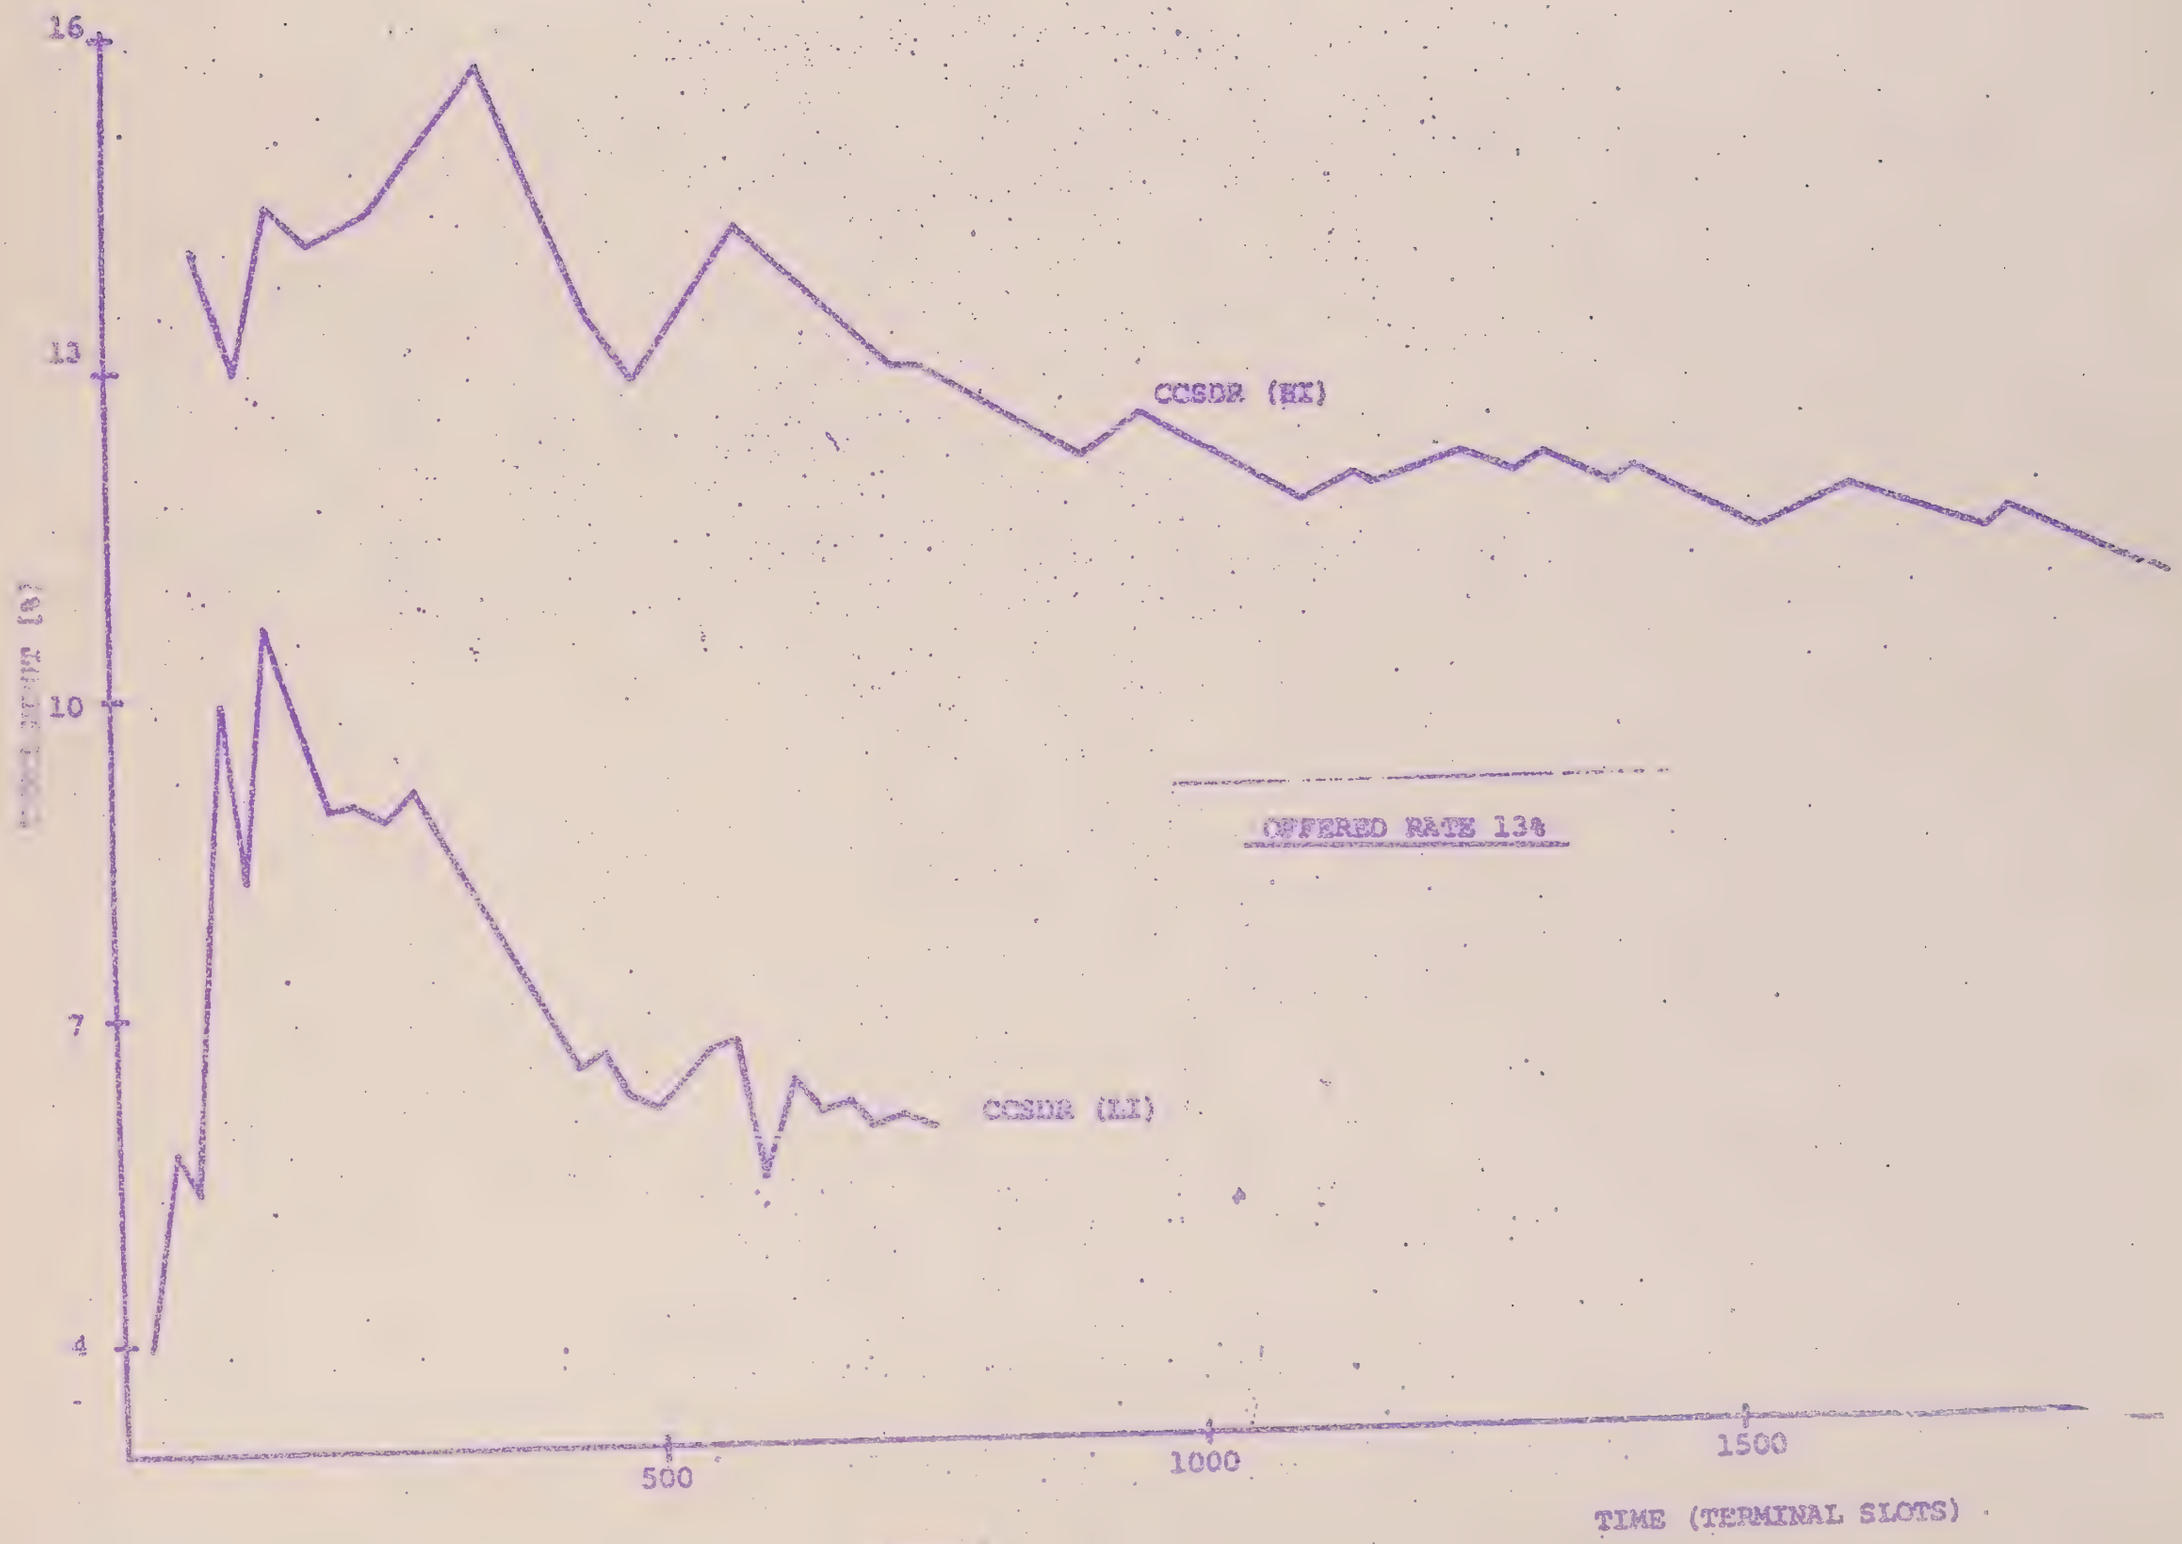
16
13
10
7
4
CCSDR (HI)
CCSDR (LI)
OFFERED RATE 13%
500
1000
1500
TIME (TERMINAL SLOTS)

	OFFERED RATE [%]	THROUGHPUT [%]	DELAY OF IP TO STATION [Terminal Slots]	RATE OF STATION RESPONSE	PROB. STATION BUSY	% OF IP BLOCKED	TOTAL % OF IP LOSS	TERMINALS REMAINING
CCSDR (LI)	13	5.95	40.11	1.14	.53	2.98	32.83	13
CCSDR (HI)	13	10.55	23.93	1.61	.43	9.83	9.83	13
CCSDR (HI) (Improved Labels)	13	12.14	16.61	2.06	.50	10.63	11.41	10

TABLE 1

labels. This experiment is discussed in detail in the next section. It is clear that the high interference system is much better than the low interference system. The only measure of the low interference system which is better is terminal blocking which is a direct result of the low interference feature. In fact, CCSDR (LI) is saturated at the offered traffic rate. This can be seen from the fact that the throughput is decreasing as a function of line, the relatively high total loss and the low station response*. The CCSDR (HI) with improved labels compared in Table 1, is better than the other two systems. This demonstrates the importance of proper labelling.

* The average number station response packets assumed for these studies is 2.0.

5. TWO VERSUS ONE DATA RATE SYSTEMS

The results of the previous section demonstrate that a better performance system is obtained when repeaters and station use high power to obtain long range despite the interference that results. We now examine the problem of whether repeaters and station should use their fixed power budgets to obtain a long range with a low data rate channel or have a short range with a high data rate channel. The following systems were studies.

- A CCSDR (HI) of the previous section with improved labels, which we denote by CCSDR. That is, we take advantage of the high range to improve the routing labels of repeaters and obtain fewer hierarchy levels. The routing labels used are shown in Figure 5, and the connectivity is shown in Figure 3.
- A Common Channel Two Data Rate (CCTDR) system with the routing labels as in Figure 1 and connectivity as in Figure 2.

In the CCTDR system, the terminal has a low data rate channel, the same rate as in the single data rate system, for communication with a repeater or station. Repeaters and station have two data rates. The high data rate is used for communication in the repeater-station network. The two data rates use the same carrier frequency so that only one can be used at a time.

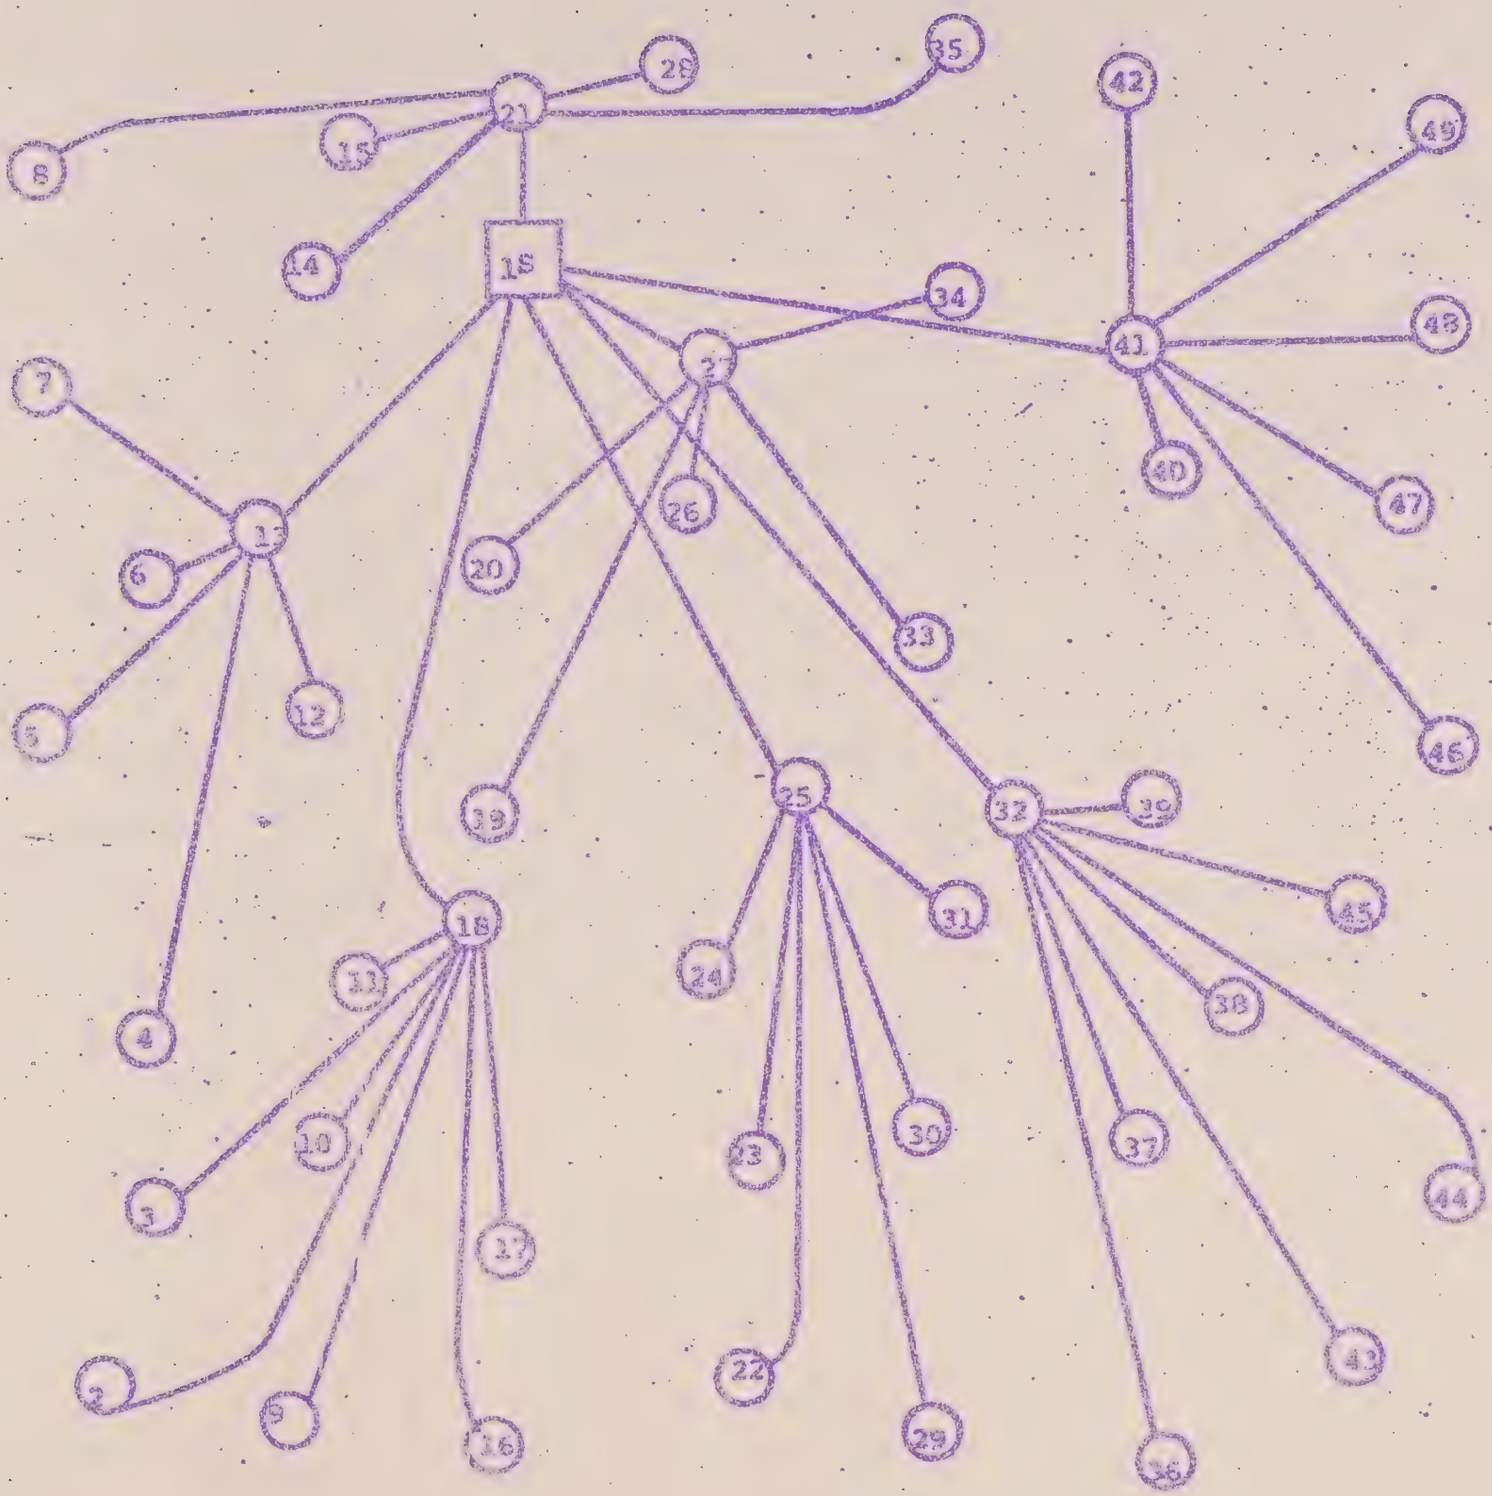
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49

FIGURE 5

The two systems are tested with offered rates of 13% and 25%.* The throughput as a function of time for the two runs are shown in Figures 6 and 7, respectively; and the summary of other measures is given in Table 2. The comparison demonstrates that the CCTRD system is superior to the CCSDR system, in terms of throughput, delay, and other measures. One can see that the CCSDR system is saturated at an offered rate of about 13%.

Effect on Blocking Level

In Table 2, one can see that one reason for the relatively low throughput of the CCSDR system at an offered rate of 25% is due to blocking. Furthermore, the fraction of time that the station is busy has decreased. This may suggest that the station may be able to handle more terminals providing they are able to enter the system. To examine this point, we ran the CCSDR system with offered rate of 25%, and relaxed the constraint for entering the system. Rather than resulting in better performance, this step resulted in reduction in blocking and increase in delay. The throughput increased to 12.63%, the blocking decreased to 18.35% and the total loss decreased to 30.73%. On the other hand, the delay increased to 57.82, the fraction of time the station is busy increased to .57, and the rate of station response decreased to 1.32.

* In the simulation runs we used the inverse square low for the relation between data rate and distance, rather than the result in [5]; this however, favors CCSDR.

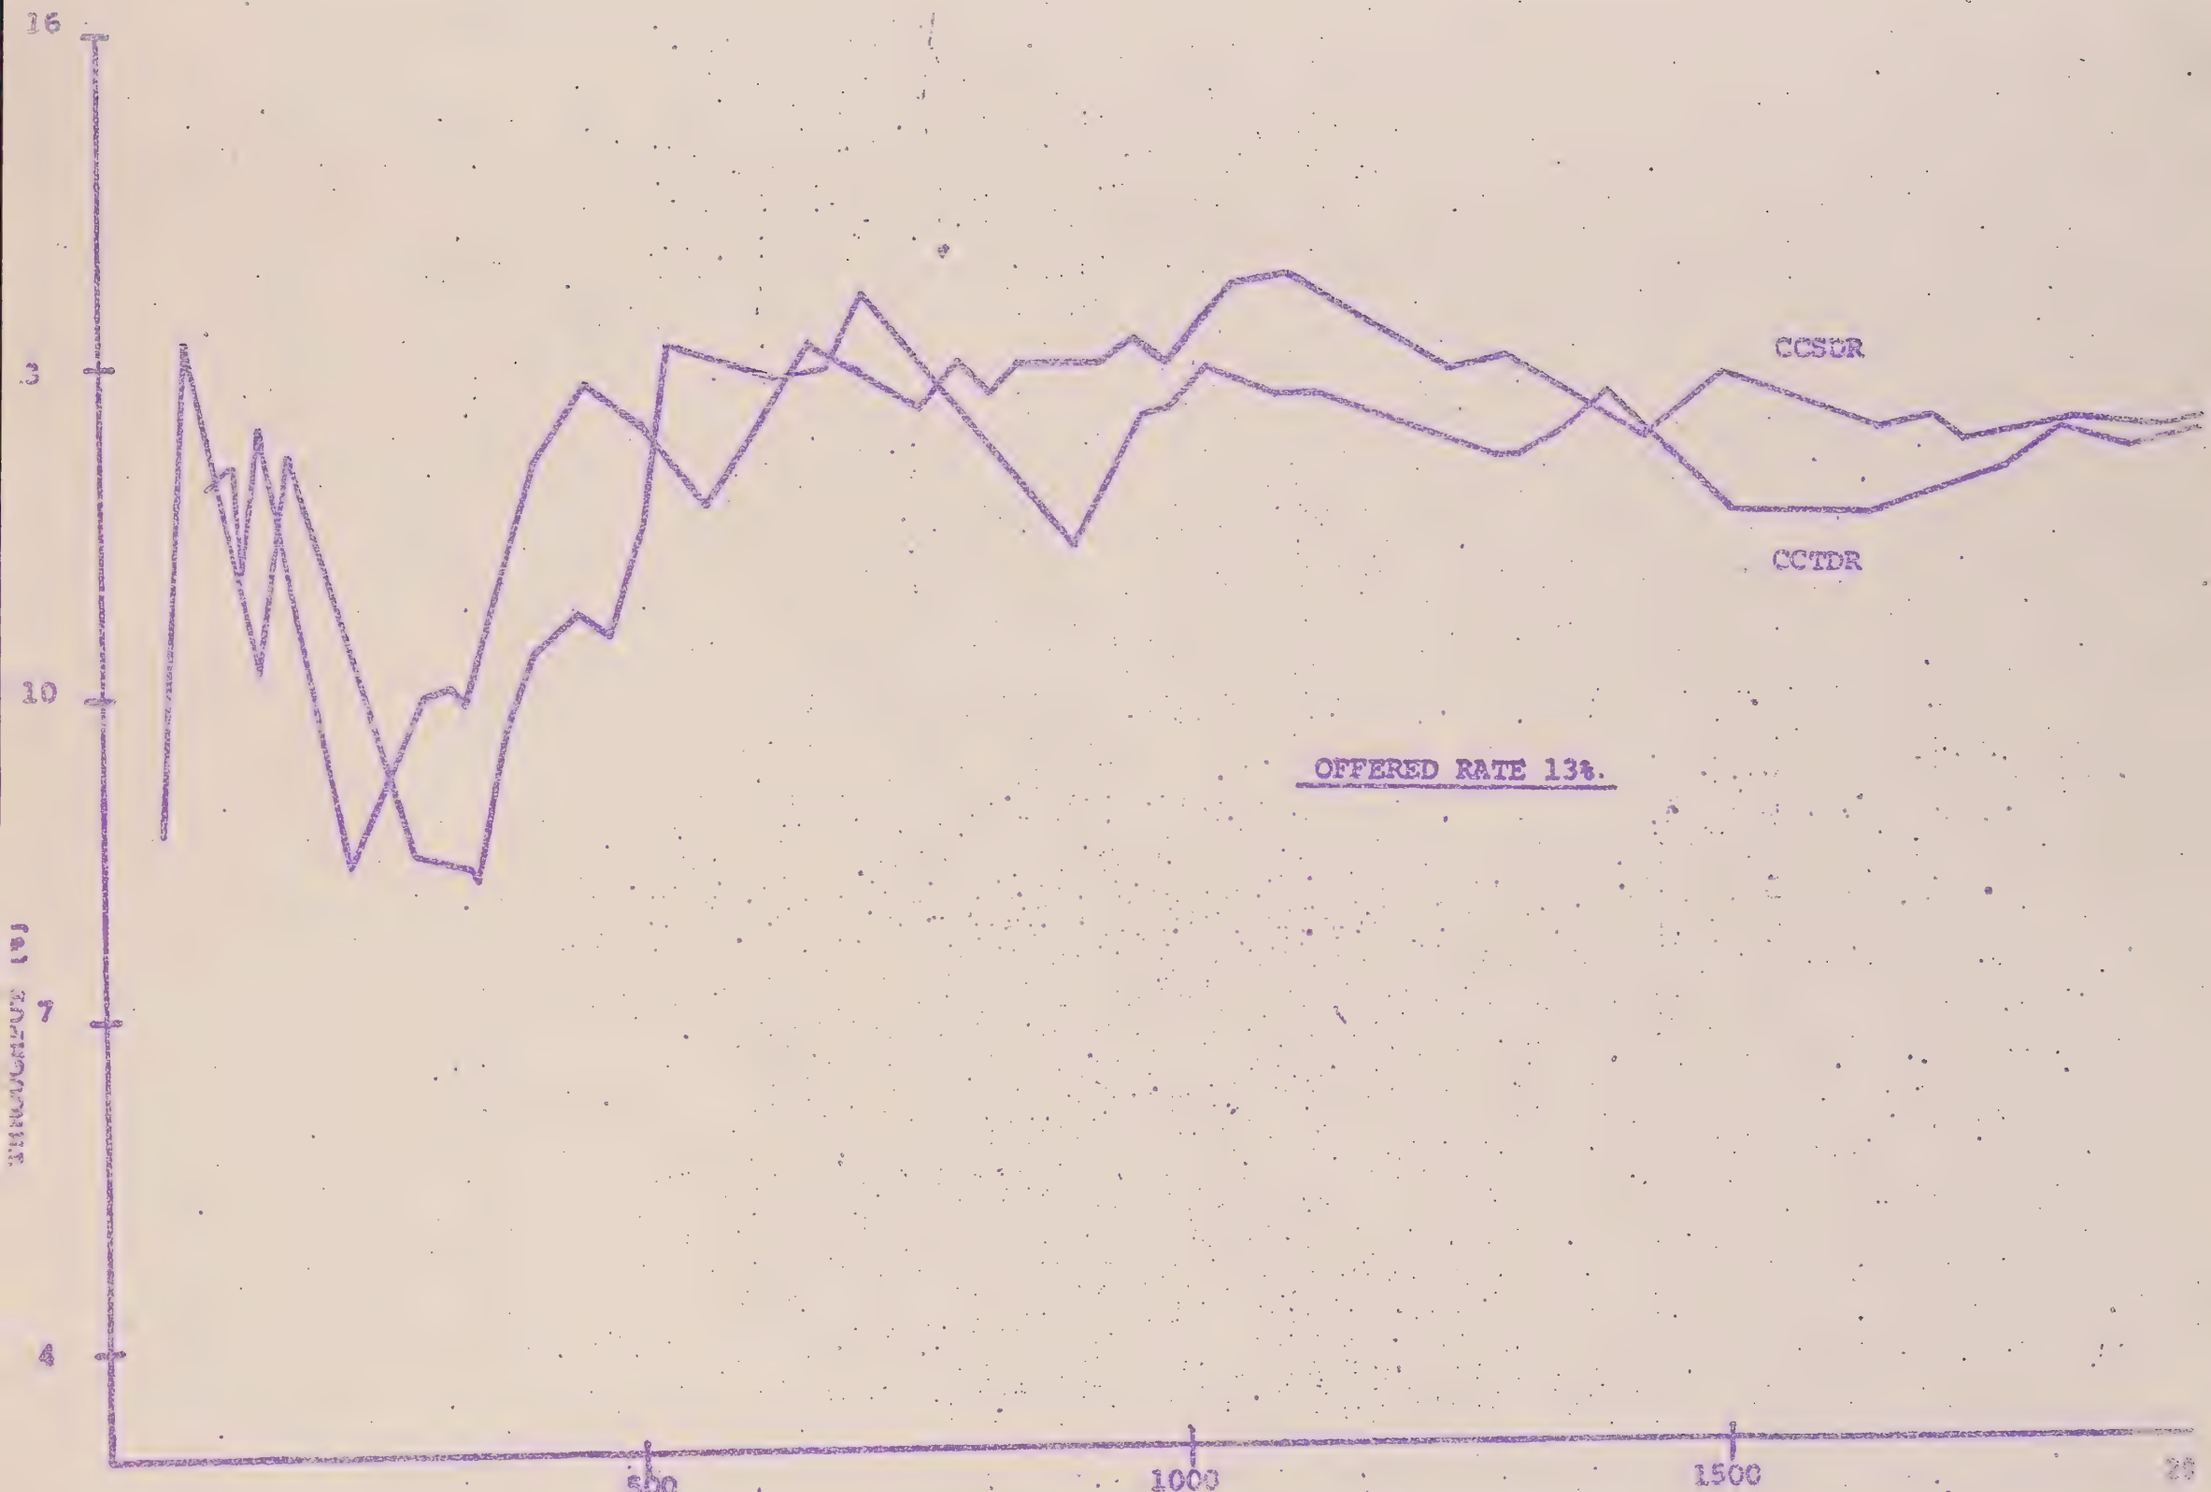
16
13
10
7
4
THROUGHPUT (%)
CCSDR
CCTDR
OFFERED RATE 13%.
500
1000
1500
2000
TIME (TERMINAL SLOTS)

FIGURE 6

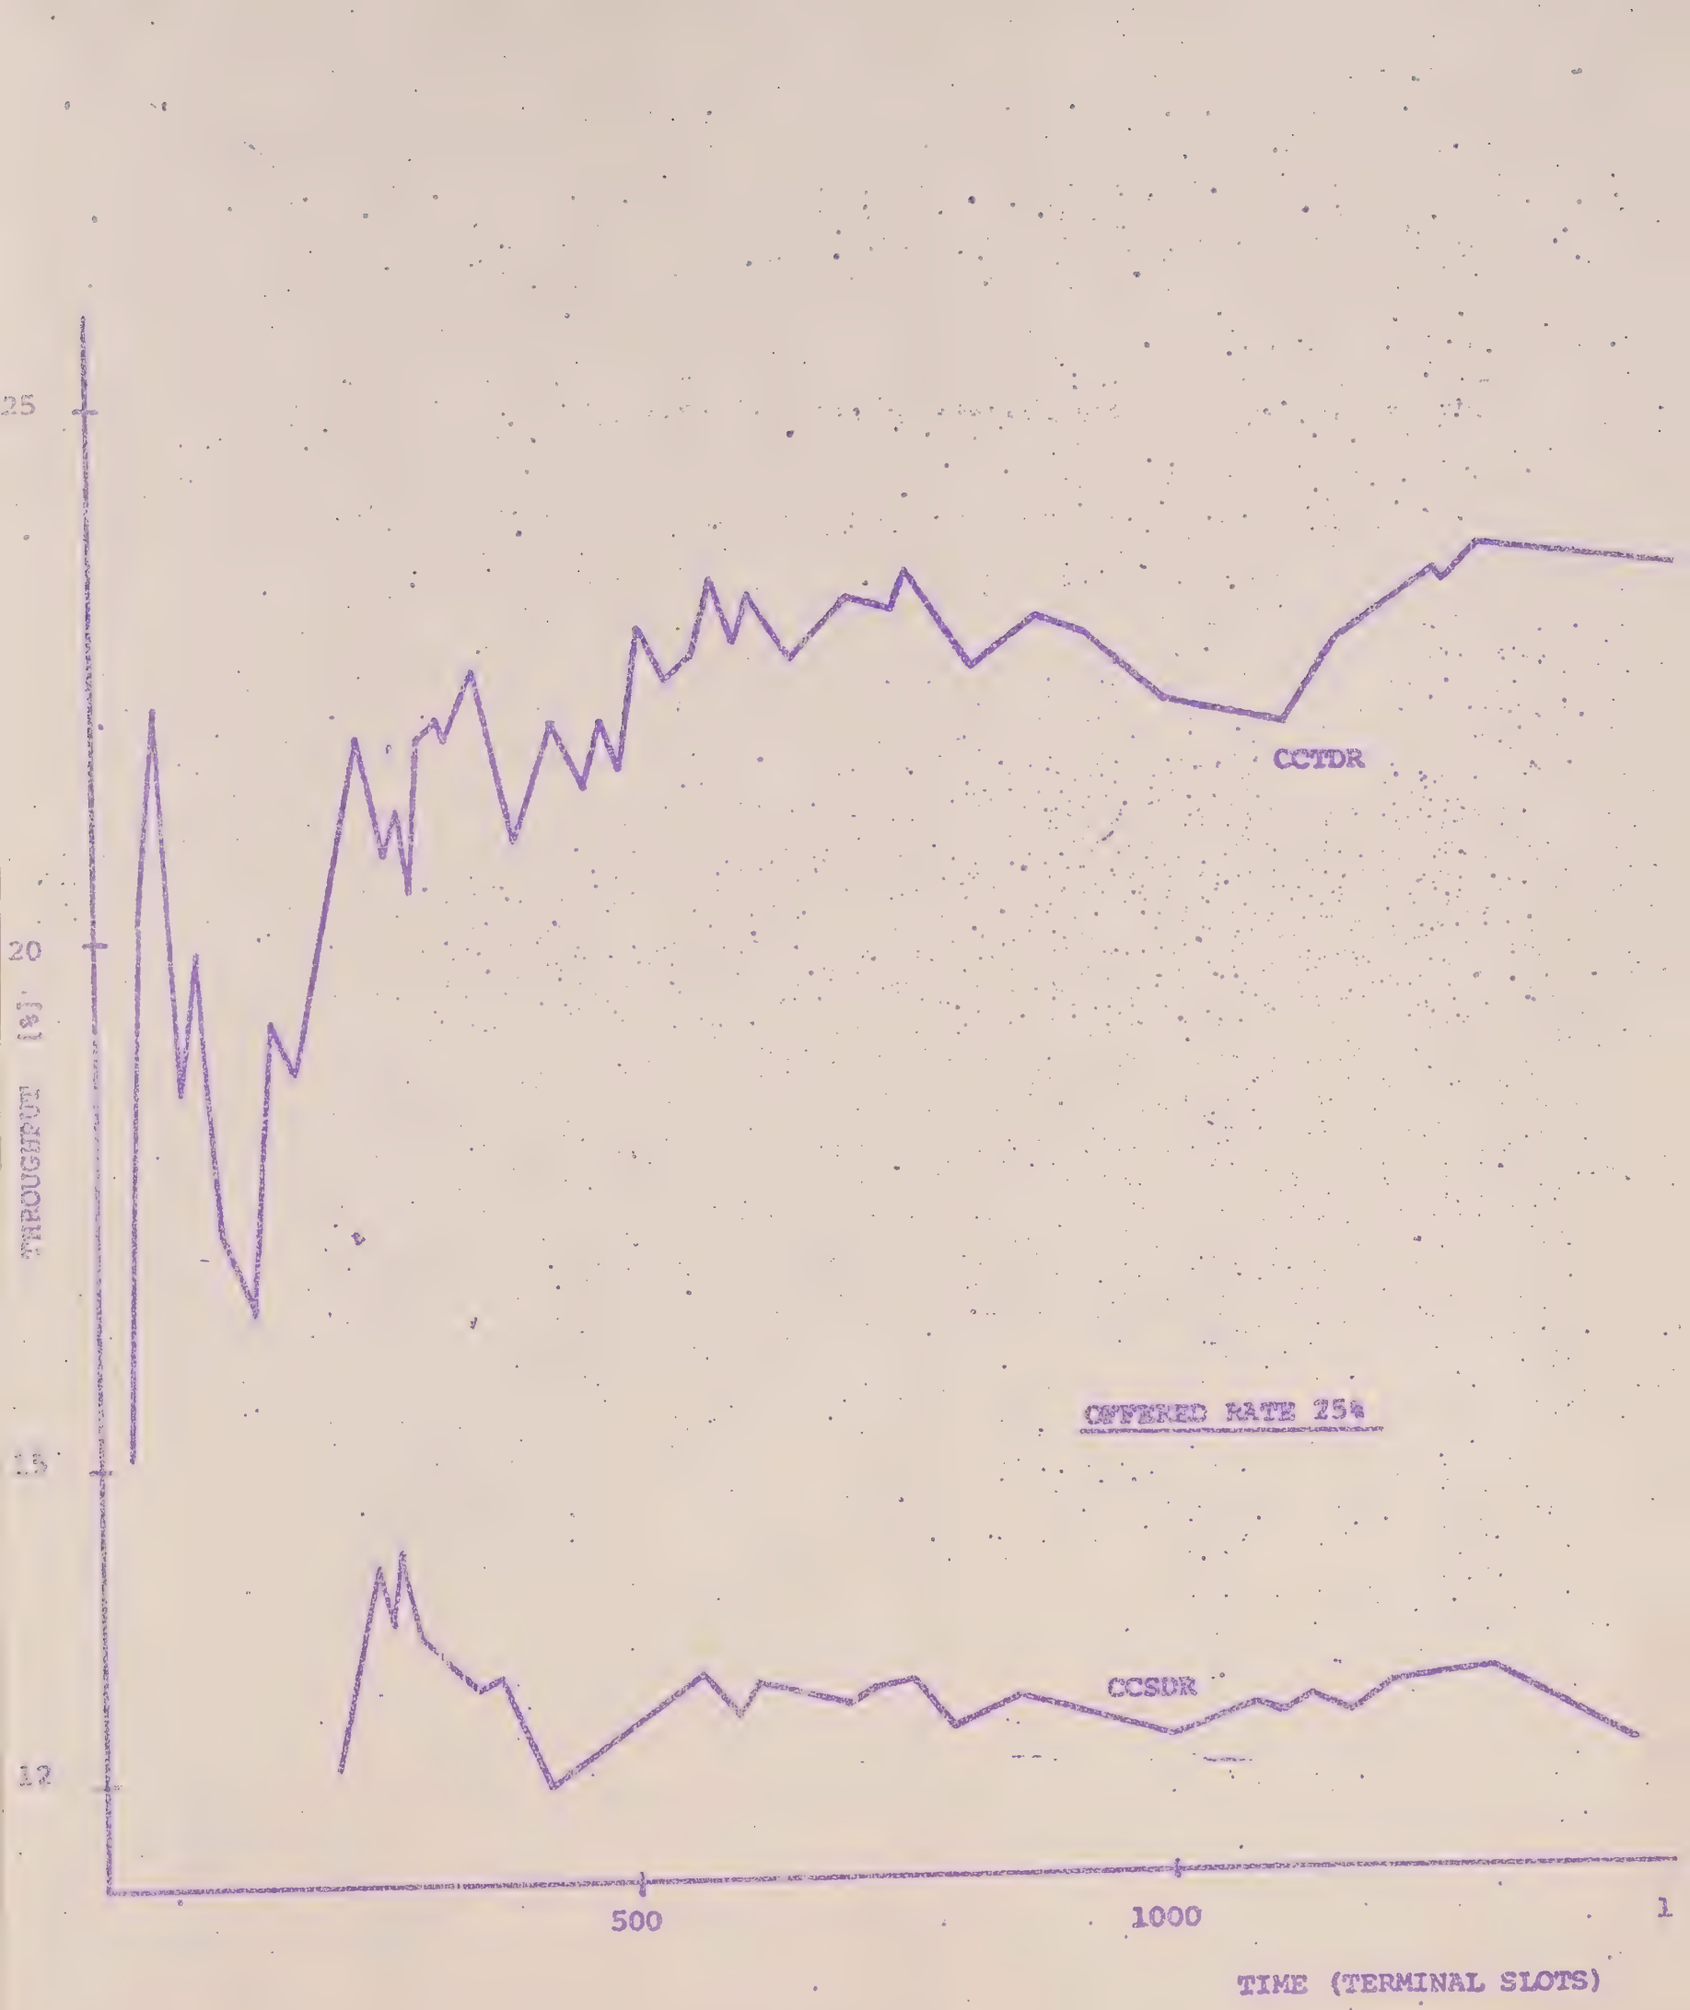
25
20
15
12
THROUGHPUT (S)
CCTDR
OFFERED RATE 25%
CCSDR
500
1000
1
TIME (TERMINAL SLOTS)

FIGURE 7

	OFFERED RATE (%)	THROUGHPUT (%	DELAY OF IP TO STATION (Terminal Slots)	RATE OF STATION RESPONCE	PROB. STATION BUSY	% OF IP BLOCKED	TOTAL % OF IP LOSS	TERMINALS REMAINING
XCSDR	13	12.14	16.61	2.06	.50	10.63	11.41	10
	25	12.20	34.97	1.61	.48	29.50	32.95	23
CCTDR	13	12.39	4.91	1.99	.26	1.59	1.59	9
	25	23.33	11.51	1.97	.31	3.31	3.31	34

TABLE 2

6. MAXIMUM THROUGHPUT, LOSS, AND DELAY OF CCSDR AND CCTDR SYSTEMS

In the packet radio system there is an absolute maximum throughput (independent of loss and delay) because of the interference characteristics. Similar to curves of throughput versus channel traffic, when the relation is known analytically [4], we draw the curves of system throughput vs. offered rate to estimating the maximum throughput. Figure 8 shows the throughput versus offered rate for CCSDR and CCTDR systems. The curves are linear for low offered rates and saturate when the offered rate increases.

For the CCSDR system one can see that the throughput is practically the same when the offered rate is increased from 13% to 25%. This and the other measures (see Table 2), (for example, the rate of station response) show that the system is overloaded at a 25% offered rate. On the other hand, the system seems to operate at steady state at an offered rate of 13% (rate of station response 2.06). A rough estimate of maximum throughput for this system would be between 12% and 15%. Similar observations of the performance measures lead to an "estimate" of between 27% and 30% for the maximum throughput of the CCTDR system.

The average delay of the first Information Packet from terminal to station, and the Total Loss, as a function of offered rate are shown in Figure 9.

There are many parameters in the simulation program which we have not experimented with (or tried to optimize) and which affect the quantities discussed above. One

parameter which is significant in determining the maximum throughput is the average number of response packets from station to terminal. The affect of this parameter has been analyzed in [6] for a slotted ALOHA random access mode. It has been shown that the maximum throughput is increased in the Common Channel system when the rate of response increases, and the maximum throughput tends to 100% of the data rate when the rate of response tends to infinity. We expect that this parameter has a similar effect for the mode of access simulated. In the results reported here the rate of response is 2.0 which is small compared to usual estimates for terminals interacting with computers. Furthermore, the relatively short terminal interaction increases the traffic overhead of the search procedure, per information packet.

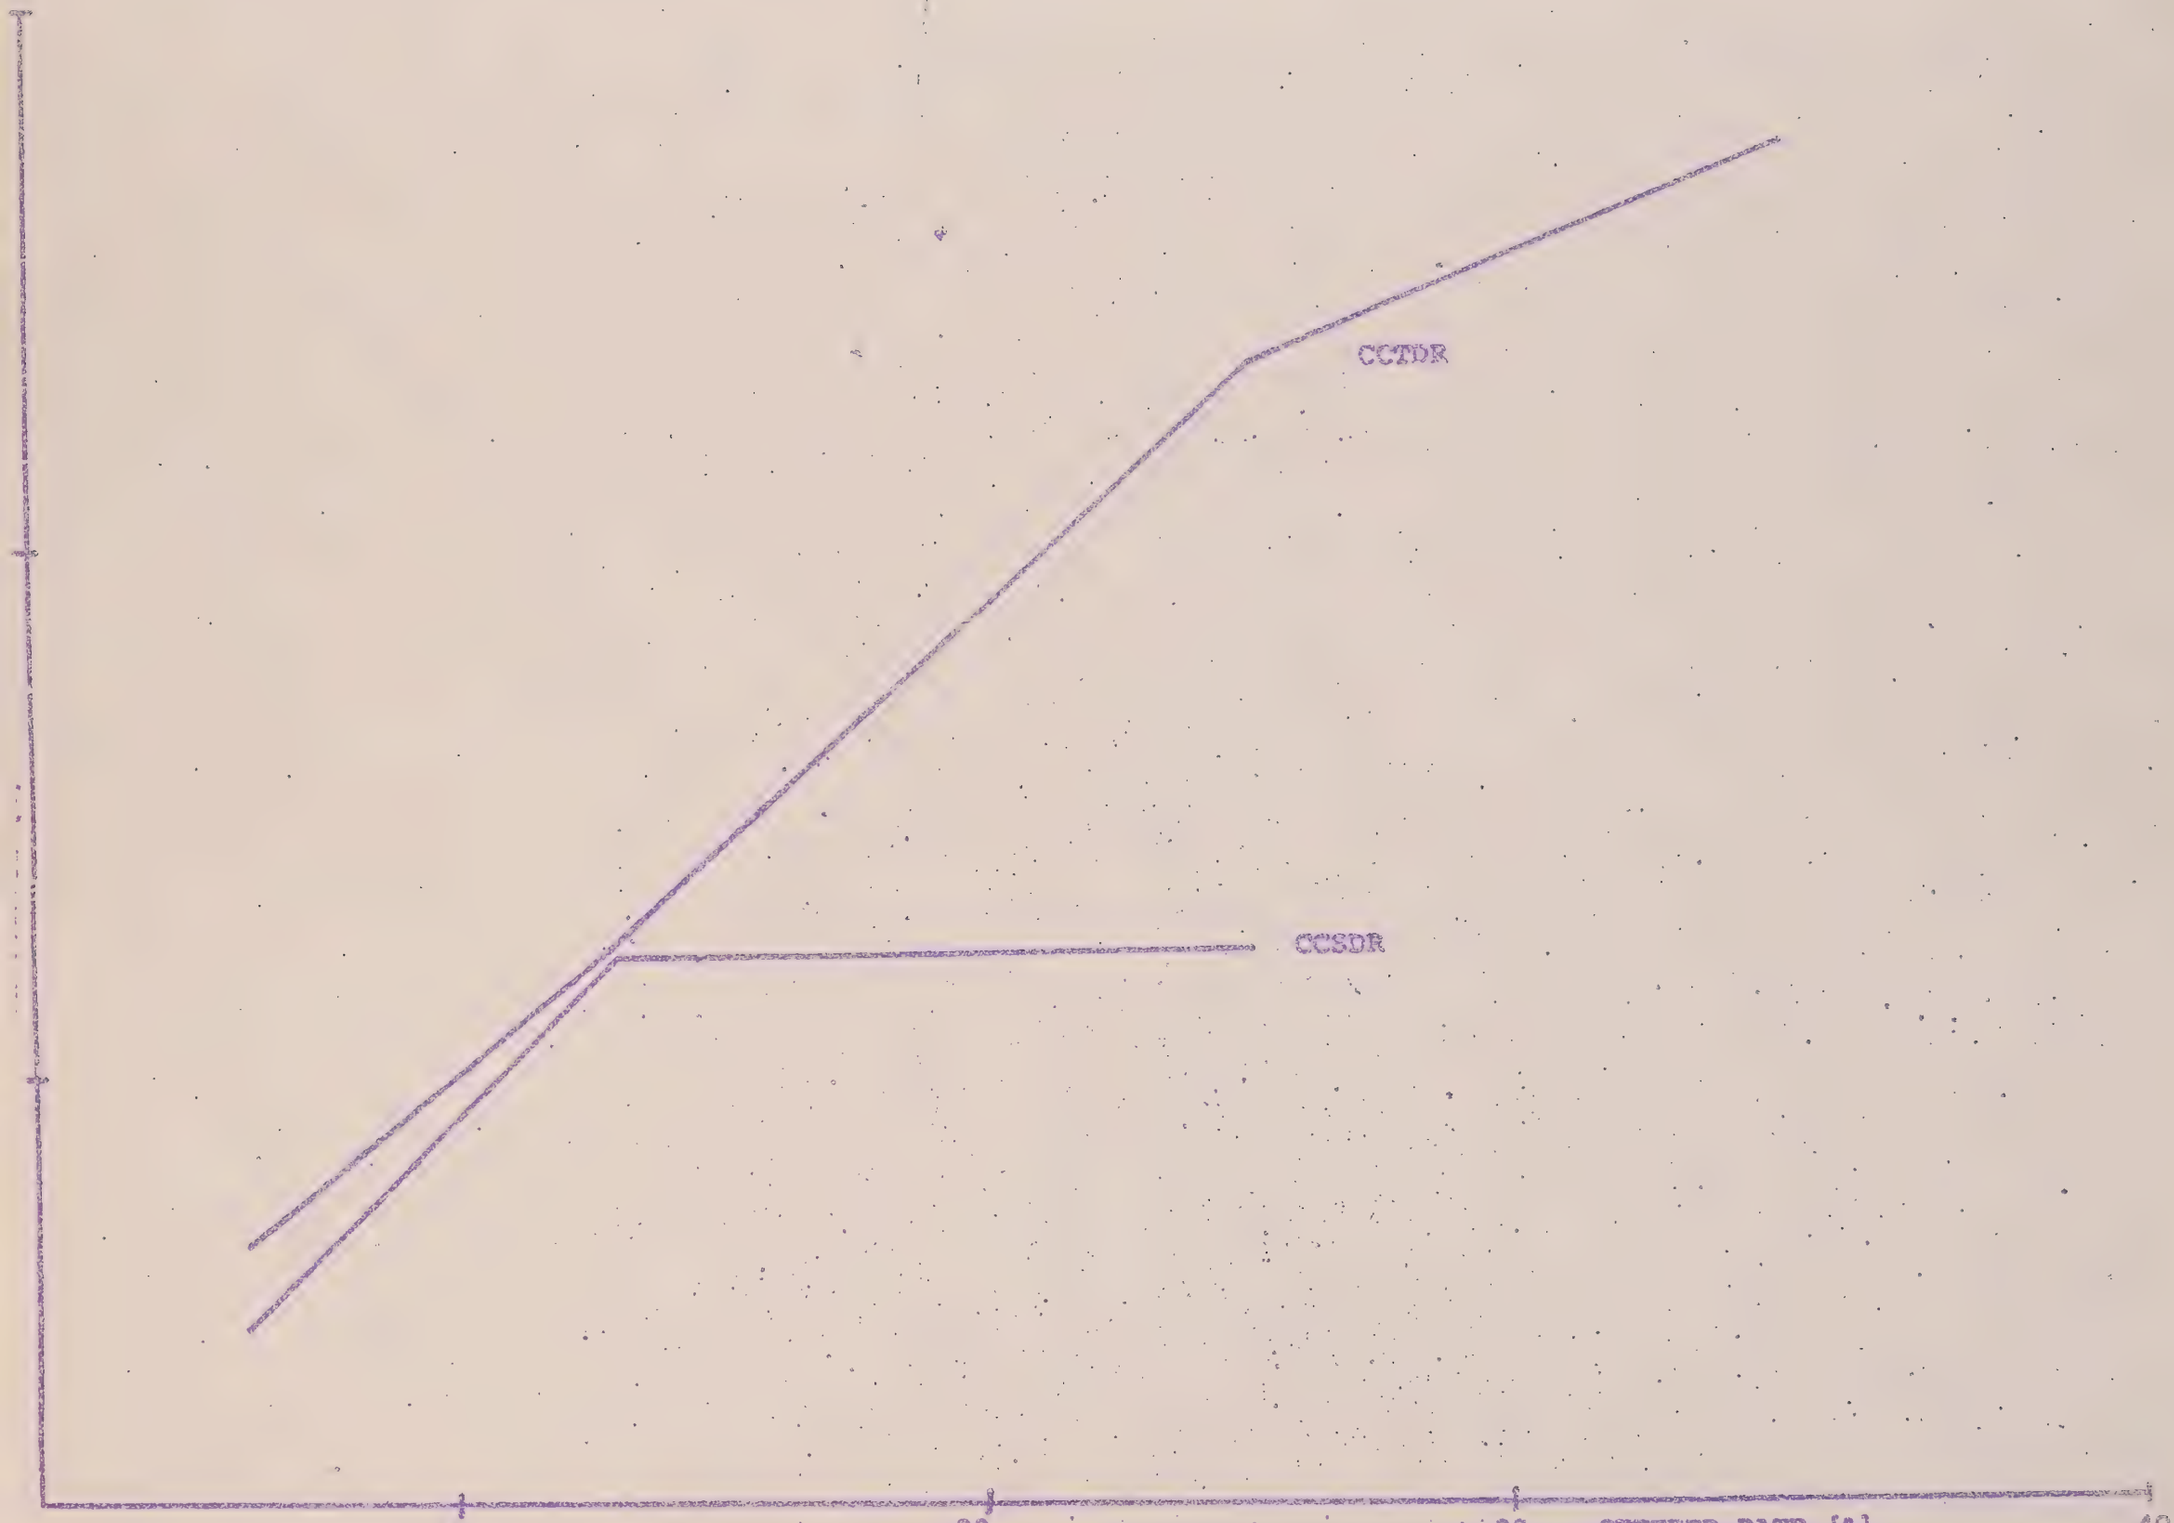

FIGURE 8

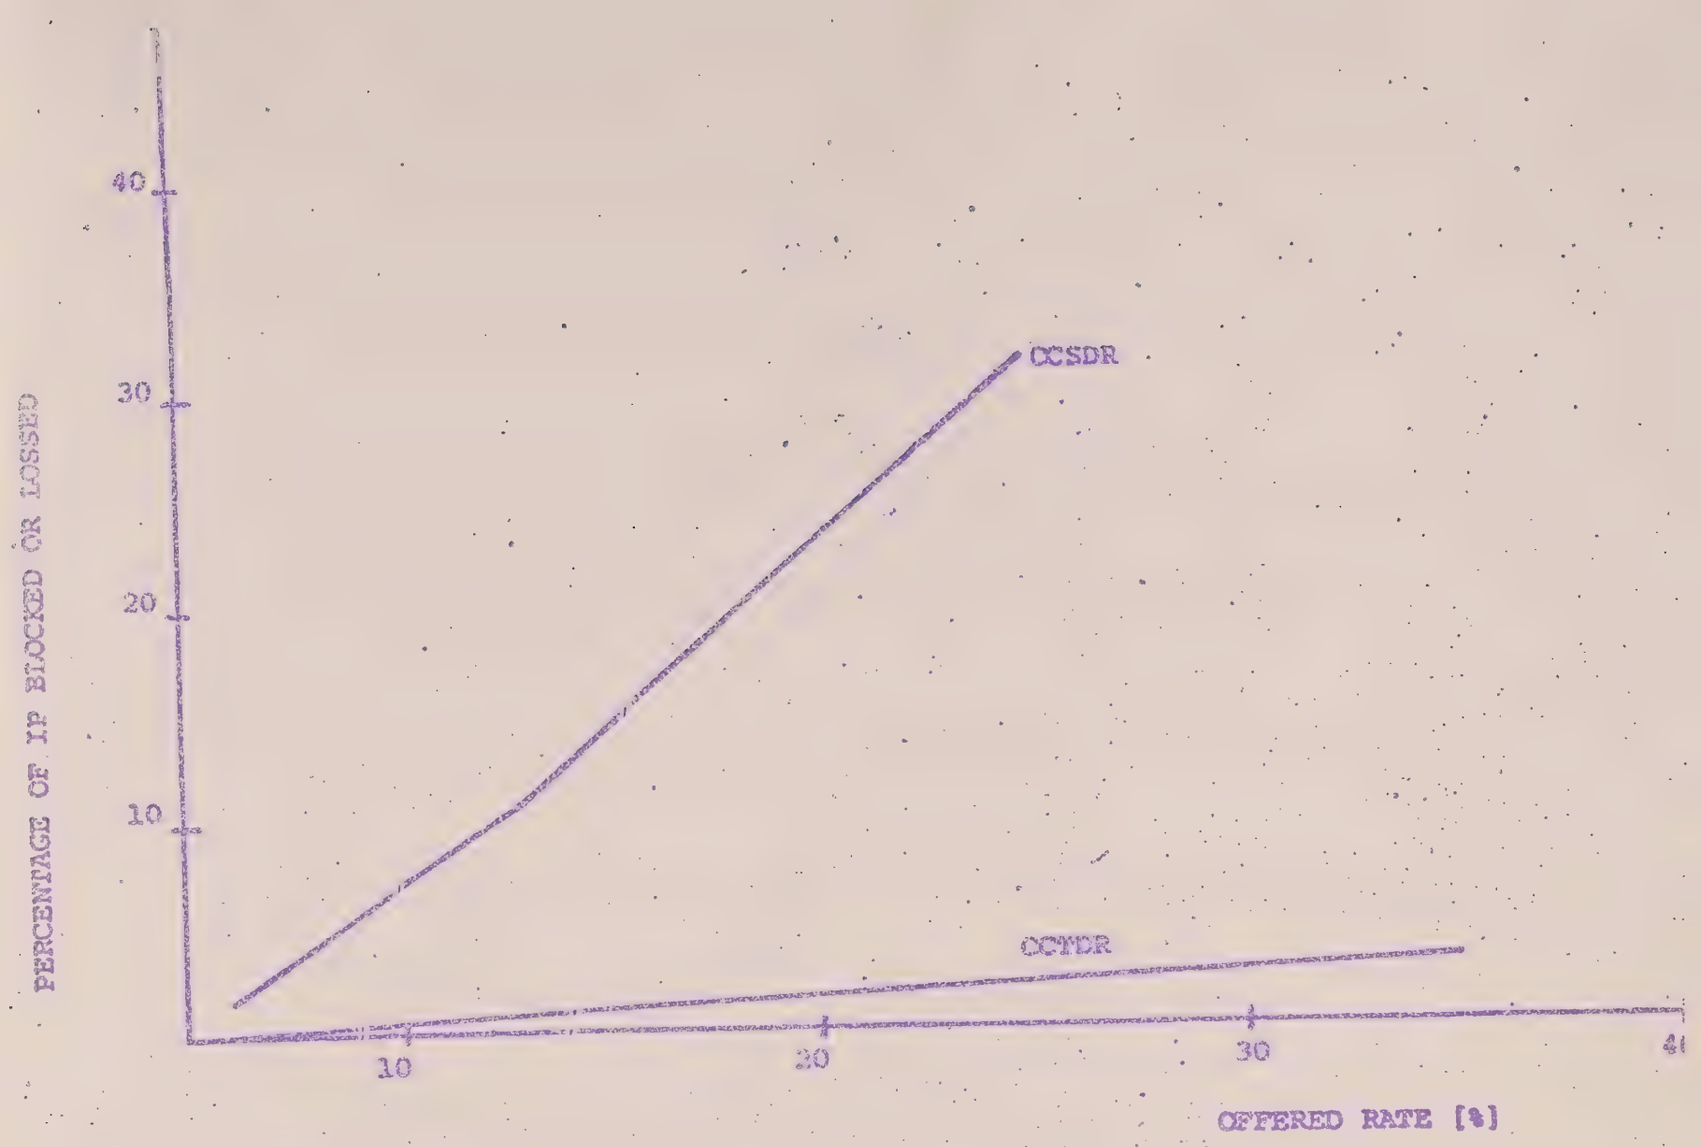
PERCENTAGE OF IP BLOCKED OR LOSSED
40
30
20
10
CCSDR
CCTDR
10
20
30
OFFERED RATE [%]

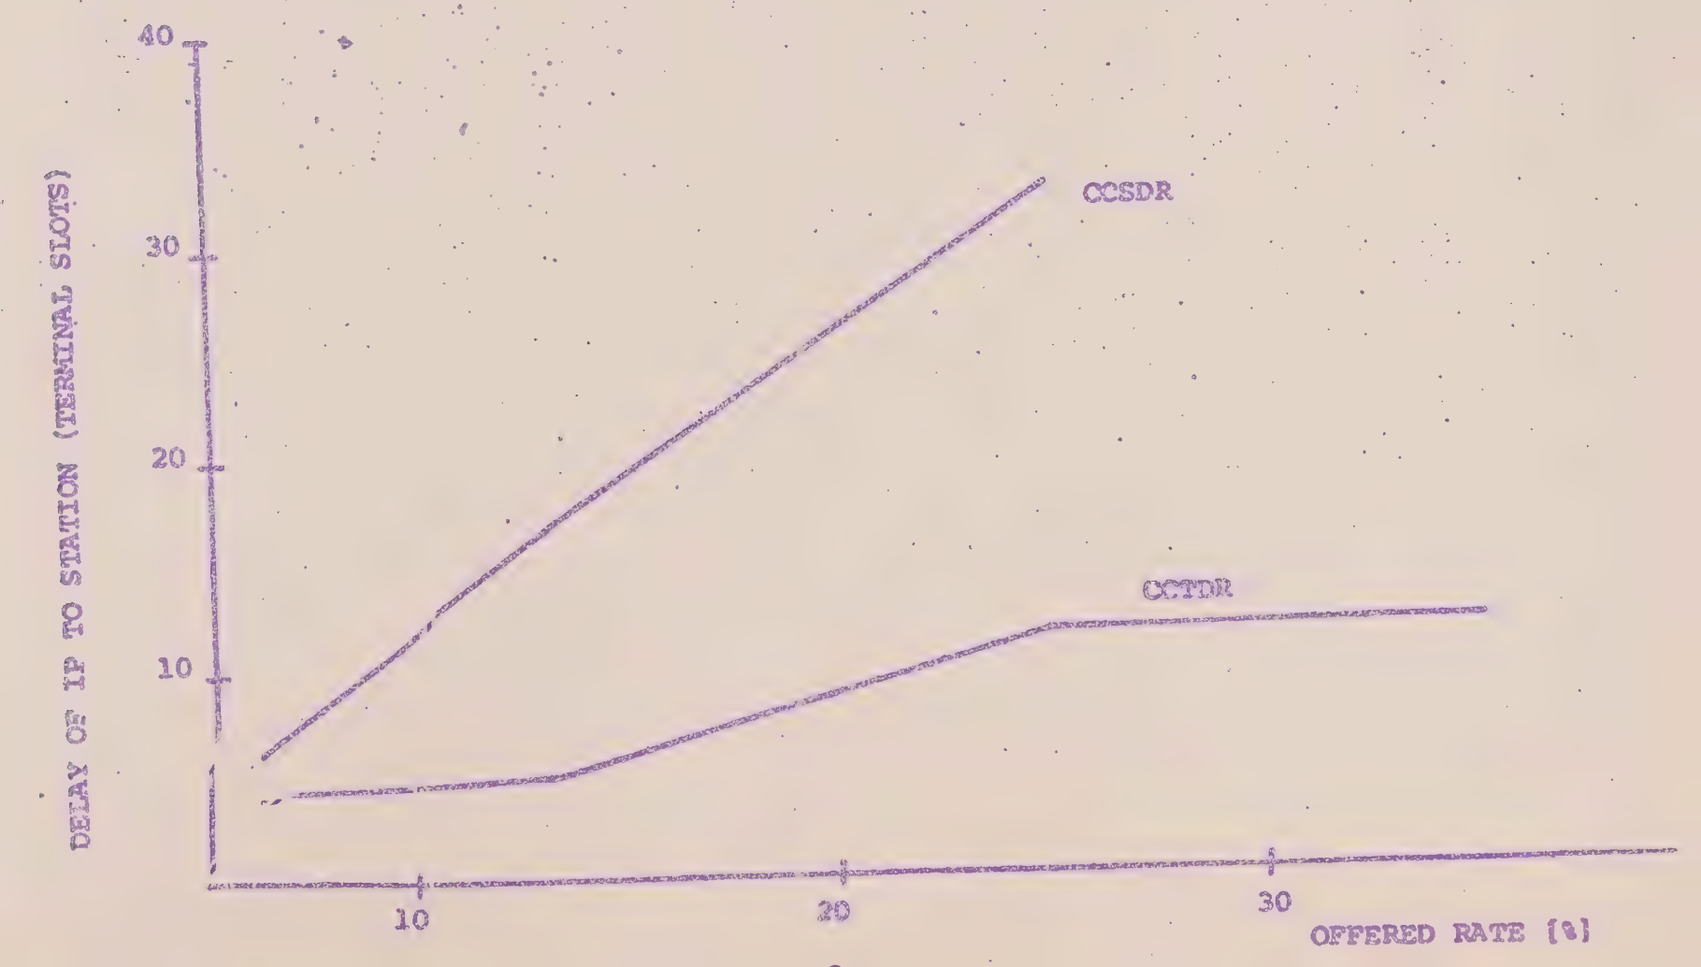
DELAY OF IP TO STATION (TERMINAL SLOTS)
40
30
20
10
CCSDR
CCTDR
10
20
30
OFFERED RATE [%]

REFERENCES

1. Network Analysis Corp: "Packet Radio System-Network Considerations," *ARPANET Network Information Center*, Stanford Research Institute.

2. Network Analysis Corp: "A Simulation of the Packet Radio Network," *ARPANET Network Information Center*, Stanford Research Institute, 1974 (b).

3. Network Analysis Corp: "Routing and Acknowledgement Schemes for the Packet Radio System," *ARPANET Network Information Center*, Stanford Research Institute.

4. Kleinrock,L. and F. Tobagi, "Carrier Sense Multiple Access (CSMA);" *ARPANET Network Information Center*,Stanford Research Institute.

5. Fralick, Stanley C., "R.F. Channel Capacity Considerations," *ARPANET Network Information Center*, Stanford Research Institute.

6. Gitman, I., R.M. VanSlyke, H. Frank: "On Presented at the Splitting Random Accessed Broadcast Communication Channels". Seventh Hawaii International Conference on System Sciences, January 8-10, 1974.

APPENDIX A

A.1 General System Description

Devices: The system simulated contains 48 repeaters and one station. These devices are assigned to corner points in a grid and their exact location is randomely distributed in the neighborhood of the corner points as shown in Figure 1. Terminals originate in time according to a Poisson process. Their geographical locations on the plane are random variables with a two dimensional uniform distribution. A terminal departs from the system when unable to establish a communication path due to congestion, after completing its communications, or after partial communication if delays are high or packets are lost.

Channel and Access Mode: The communication channel is shared for transmission in both directions, to the station or from the station to terminals. The channel access mode is the Non-Persistent Carrier Sense Multiple Access (CSMA) [4]. That is, when a packet is ready for transmission, the device senses the channel and transmits the entire packet if the channel is idle. If the channel is busy, the device reschedules the packet for some future random time at which it senses the channel again, and the procedure repeated.

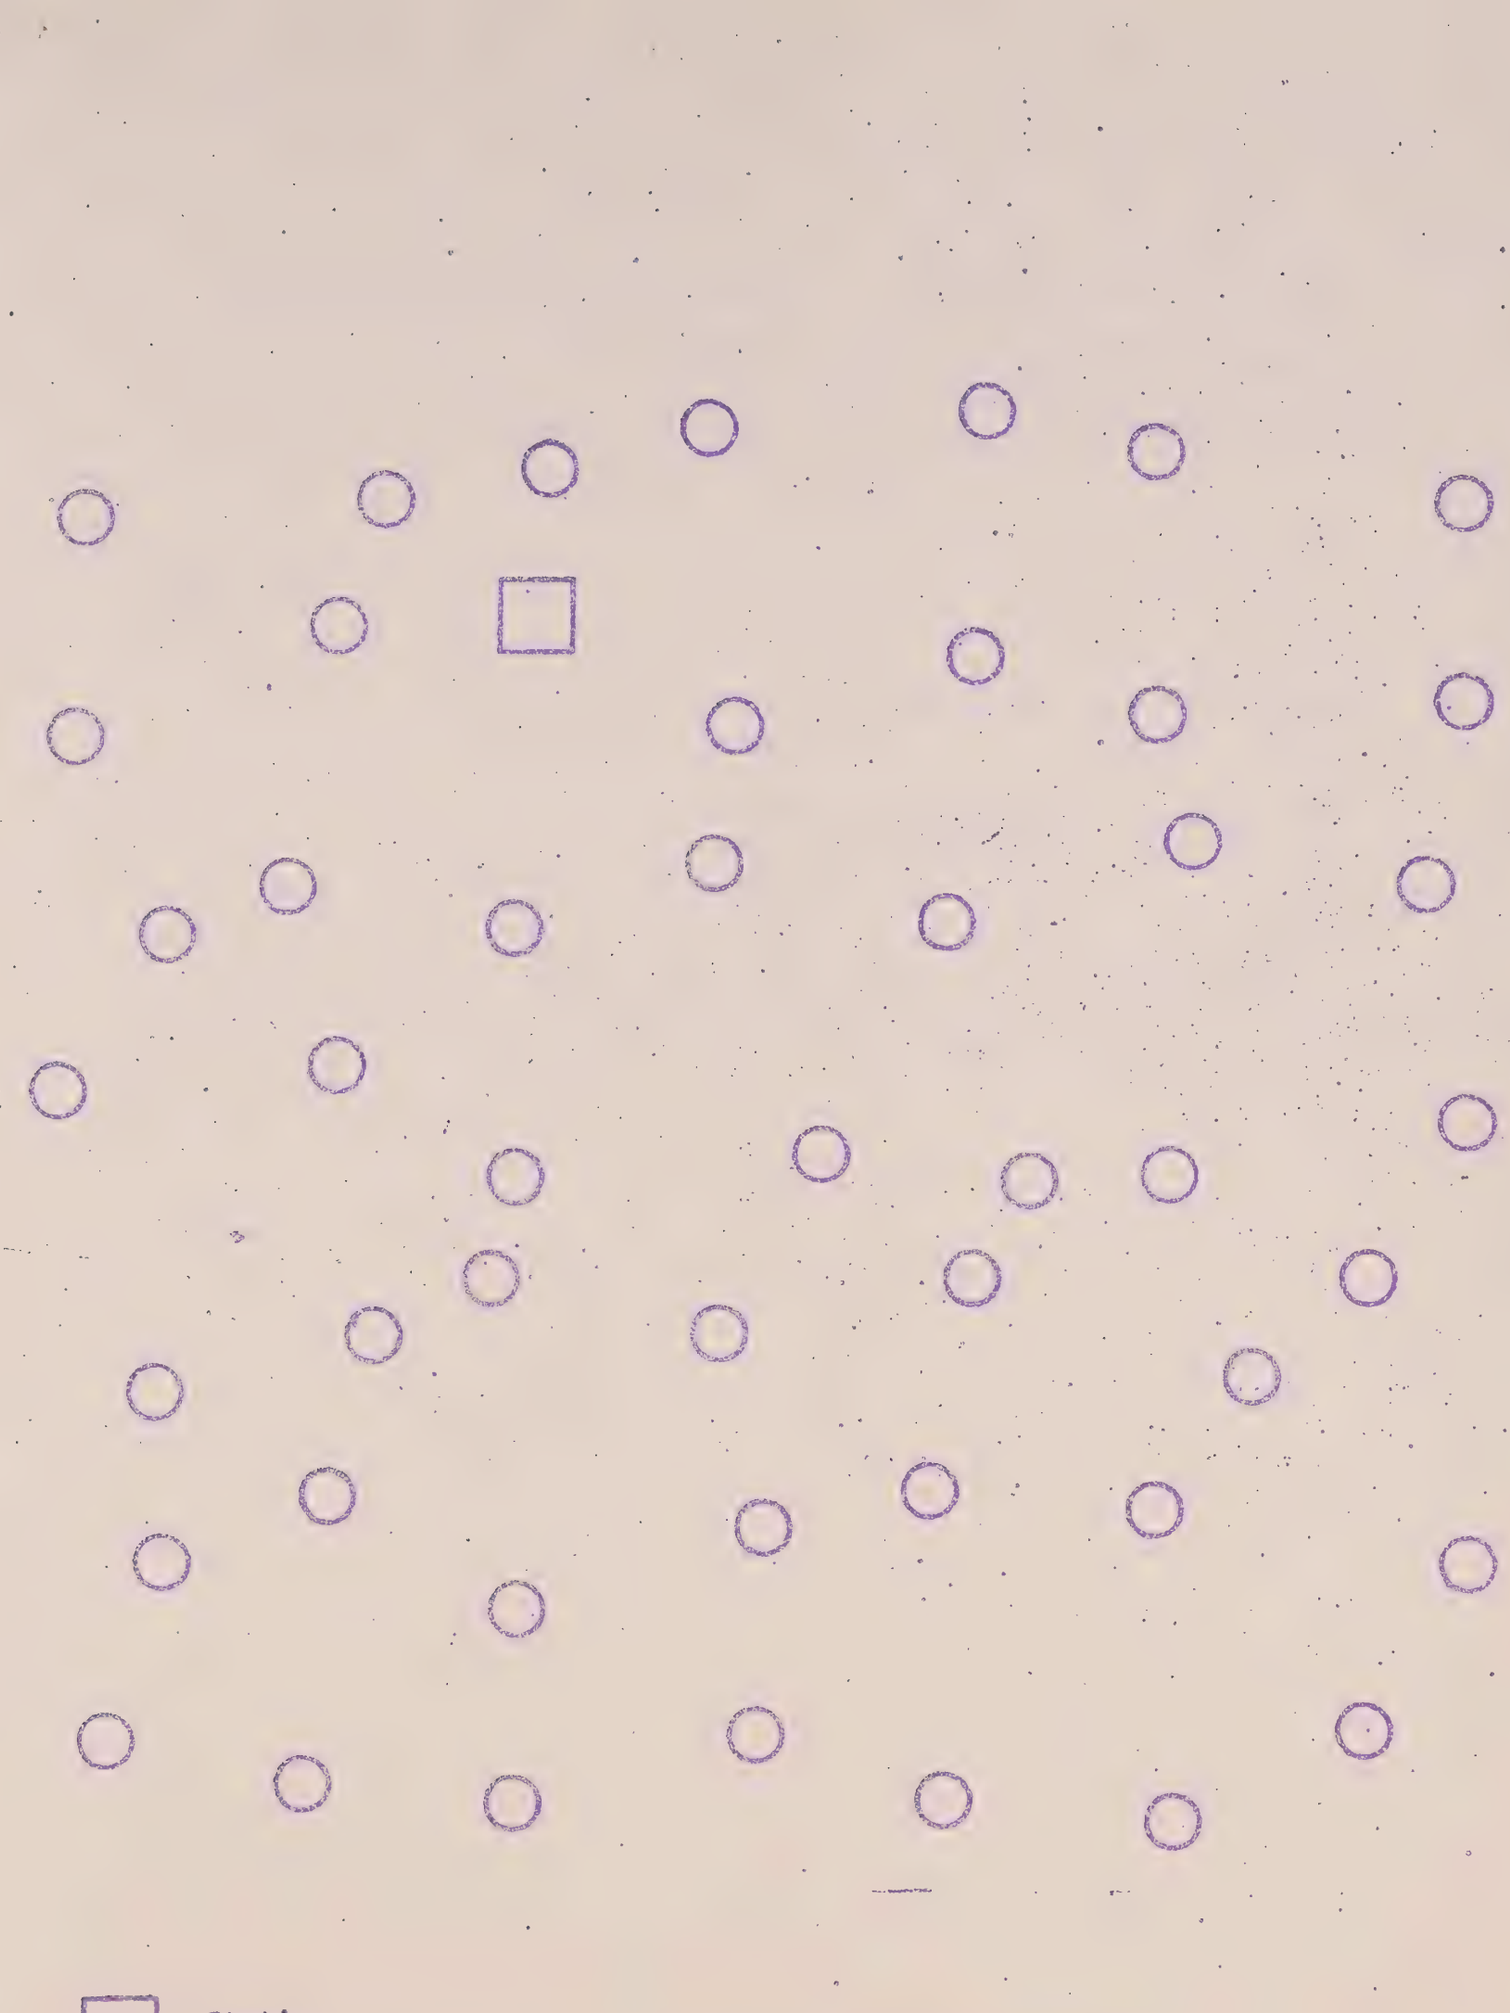

FIGURE A1: DEVICE LOCATIONS

Capture: A zero capture system is simulated. That is, whenever the reception of more than one packet overlap in time, none of the packets is correctly received.

Packet Types Simulated

IP - Information Packet

ETE - ETE Ack Short Packet (Assumed to be 10% of the length of an information packet.)

SP - Search Packet, transmitted by Terminal or Repeater to all devices and aimed to identify a specific receiver (short packet)

RSP - Response to a Search Packet, transmitted by Repeater or Station which received a SP and is available for handling packets. This packet contains the label of the transmitting device and is addressed to all devices.

A.2 Routing and Acknowledgements

A formal presentation of the routing and acknowledgement schemes is given in [3]. Devices are assigned hierarchical labels which result in directed (point-to-point) transmission in the packet radio system. Information packets are addressed to particular devices. Devices use omnidirectional antennas; thus, an information packet is received and interferes with all devices within the effective range of the transmitting device.

Two acknowledgements are used. An End-To-End Acknowledgement (ETE Ack) between terminals and station to insure message integrity and a Hop-by-Hop Acknowledgement (HBH Ack) along the communication path. The HBH Ack is a passive "Echo" Acknowledgement. That is, when device i transmits a packet to device j, it waits sufficient time to allow device j to repeat the packet to the next device; when this transmission by device j is received by i, it considers it as an acknowledgement. If device j is an end device (station or terminal) which need not repeat the packet, it specifically repeats the header for this purpose.

As discussed in reference [3], a repeater can attempt "forward" or "forward and backward" search when encountering blocking on the shortest path. In the simulation we use one SF which is equivalent to "forward and backward" search. This simplifies the repeater since it need not construct the various search packets or identify these when receiving one. Also, the search packet has the same format

whether transmitted by a repeater or by a terminal.

The Response to Search Packet (RSP) is transmitted by repeaters after a random waiting time. This time randomization is essential in a no-capture system. Otherwise, if more than one receiver wish to respond to the SP, the transmissions will overlap at the searching device. The station, on the other hand, transmits the RSP immediately after receiving the SP (Note that the channel is idle at this time since otherwise the SP would not have been correctly received). The above enables searching devices within range of an idle station to communicate directly with the station.

A repeater makes one attempt to transmit a RSP and if the channel is busy it discards it, rather than store the RSP for future transmissions. This allows control of the level of terminal blocking by specifying the number of transmissions of the SP by a terminal. Thus, when the system is conjested in the geographical neighborhood of the terminal it will not be able to "enter" the system. This feature also makes repeaters more available for handling information packets.

The Maximum Handover Number (MHN) was initially introduced into a packet header as a control to prevent packets from cycling in a loop of repeaters or being propagated to a large distance [3]. The MHN is decremented at every hop and the packet is discarded when the MHN

reaches zero. The MHN also has an important role in the HBH Echo Ack scheme and its implications.

In a point-to-point network, such as the ARPANET, the channels are fixed so that when node j receives a packet on channel k it knows the device, say i, which has transmitted the packet to it, and thus can transmit a specific HBH Ack to device i. In the Packet Radio System, however, this information is not available. Therefore, the HBH Ack must be independent of the path that the packet travels on. The HBH Echo Ack test used includes the following:

(i) Identification of the packet

(ii) Tests that the MHN of the Echo received is smaller than the MHN of the packet stored. This insures that the packet has advanced along the path, rather than being a retransmission from devices that had the packet previously.

The HBH Echo Ack has several advantages over a specific HBH Ack:

(i) It simplifies the repeater (hardware and software) so that it need not construct and manage acknowledgement packets.

(ii) It reduces the traffic overhead of transmitting specific acknowledgements. This is most significant in a broadcast network.

(iii) It enables acknowledgement of several devices at a time; in particular, all devices which store the packet with a MHN larger than that received are acknowledged.

(iv) It enables shortening the transmission path, as described below.

Since the RSP's by repeaters are randomized in time, a terminal frequently does not identify the repeater nearest the station within

range of the terminal. In fact, if two repeaters are labelled on the same path to the station and both are within range of the terminal, there is a higher probability that the terminal will identify the repeater farther away from the station since, on the average, the latter handles less traffic. Suppose a single data rate channel is used throughout the transmission path between terminal and station, the terminal identifies a packet transmitted to it by its specific terminal ID, and the station can recognize any packet destined for it. Then a communication path as shown in Figure A2 may be established. In Figure A2, the terminal is within an effective range to R4, R5, and R6; and the station within an effective range to R1, and R2. The terminal shown identified R6 as the repeater to which it transmits. The path from terminal to station will usually be: T → R6 → R5 → R4 → R3 → R2 → S; and from the station to the terminal S → R2 → R3 → R4 → T. The end devices, terminal and station transmit the Echo Ack immediately after receiving the packet, and transmit it with MHN=0; thus they acknowledge all devices which still store the packet. In particular, when R4 transmits a packet towards the terminal it is addressed to R5, however, the terminal may receive this packet and acknowledge both R4 and R5 simultaneously. Similarly, when R2 transmits towards the station it addresses the packets to R1, when the packet is received by the station it acknowledges both R1 and R2 simultaneously.

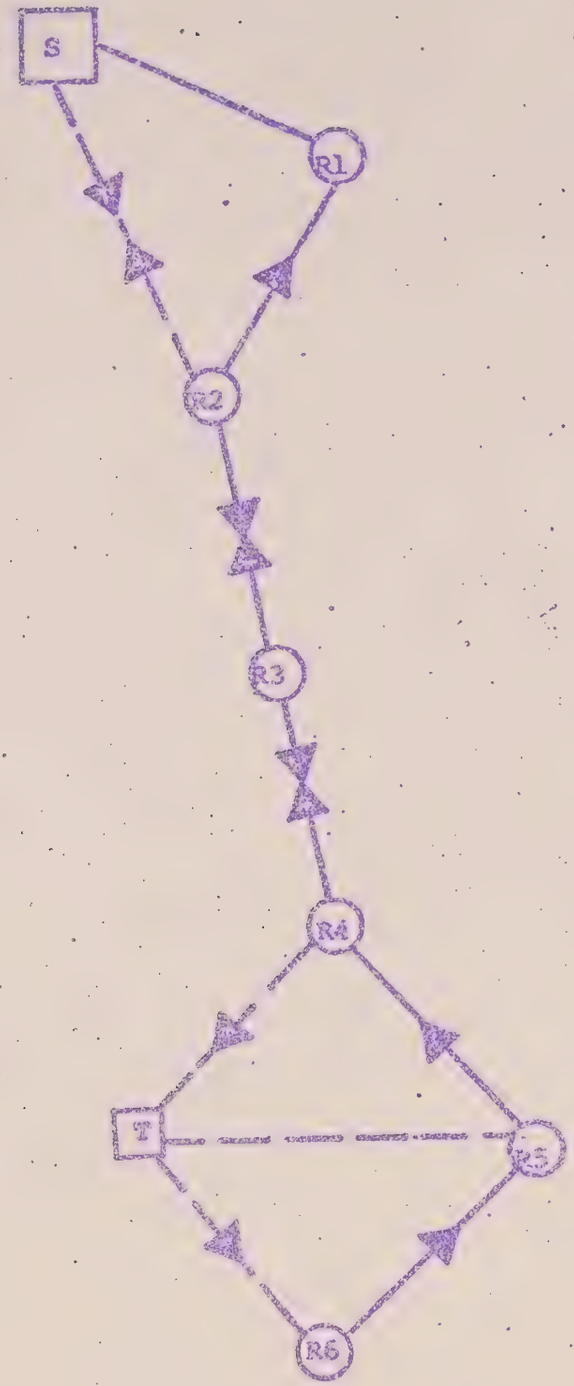

FIGURE A2: The solid lines indicate the labelled path between the repeaters and station. The dashed lines indicate the effective connectivity of terminal and station of repeaters. The arrows indicate the practical transmission path for the particular terminal, to and from the station.

A.3 Logical Operation of Devices

Terminal: When a terminal originates a message, it begins to transmit and retransmit a SP to identify a specific receiver. If it does not identify a specific receiver after a specified number of transmissions, it departs from the system. We say that such a terminal is blocked. When a terminal identifies a specific receiver, it substitutes the label and MHN sent by the receiver into its IP and begins to transmit its IP. The IP is retransmitted after short waiting periods of time until a HBH Echo Ack is received. At that time the terminal stores the IP for a longer period of waiting after which the IP is reactivated if an ETE Ack is not received.* The terminal is expecting several IP's from the station, each of which is ETE acknowledged by the terminal. When all IP's from the station are received and ETE acknowledged, the terminal departs from the system.

---

* We use the term retransmission when a device waits a relatively short period of time (less than 2 IP slots) and is awaiting a HBH acknowledgement. We say that a packet is reactivated when an end device stores the packet, awaiting an ETE Ack. When a packet is reactivated it goes through the whole process of retransmissions.

Repeater: A repeater does not distinguish between IPs or ETE Acks, except for their transmission time. When an IP (or ETE Ack) is received by a repeater (addressed to it) and the repeater has available storage, it stores the packet, decrements the MHN, modifies the packet label according to the routing and begins to transmit and retransmit the packet, awaiting an Echo Ack. When an Echo Ack is received the repeater discards the packet. When a repeater is not successful in transmitting along the "shortest" path it begins to search for an alternate receiver by transmitting SPs. When one is found it transmits the entire packet to it; otherwise, it discards the packet. When a repeater receives an SP it checks whether it has available storage; if it does, it makes one attempt to transmit a RSP and then discards it. When a repeater receives a RSP it tests whether it needs one; if it does, it uses the label, otherwise it discards it. The repeater currently simulated has buffer storage for two packets; one exclussively used for packets directed towards the station and the second for packets toward the terminal. In addition, the repeater can inspect all packets that it receives, which are stored in common arrays in the simulation program . Thus, from a practical viewpoint, buffer storage for three packets per repeater are provided in the simulation program.

Station: The storage organization in the simulated station is shown in Figure A3.* There are two queues for active packets. Packets in these queues are active in the sense that they are retransmitted after short random periods of waiting until an Echo Ack is received. The active queue for long packets contains IPs from the station to terminals. Once an IP is Echo acknowledged it is stored in the passive queue for a longer period after which it is reactivated if an ETE Ack from the terminal is not received. The active queue for short packets contains ETE Acknowledgement packets to terminals, and these have priority over the long active packets. The ETE Ack packets are, obviously, discarded once an Echo Ack is received. The point-to-point (PTP) network queue simulates the interaction of the packet radio network with a PTP network. When a new IP is received from a terminal, it is stored in the PTP network queue for a random time, after which a response message containing several response IPs to that terminal are generated and placed into the active queue for long packets. The station responds immediately to SPs, and ignores RSPs.

* Figure 3 shows only the storage of IP's and ETE Acks for transmission to terminals in the packet radio network.

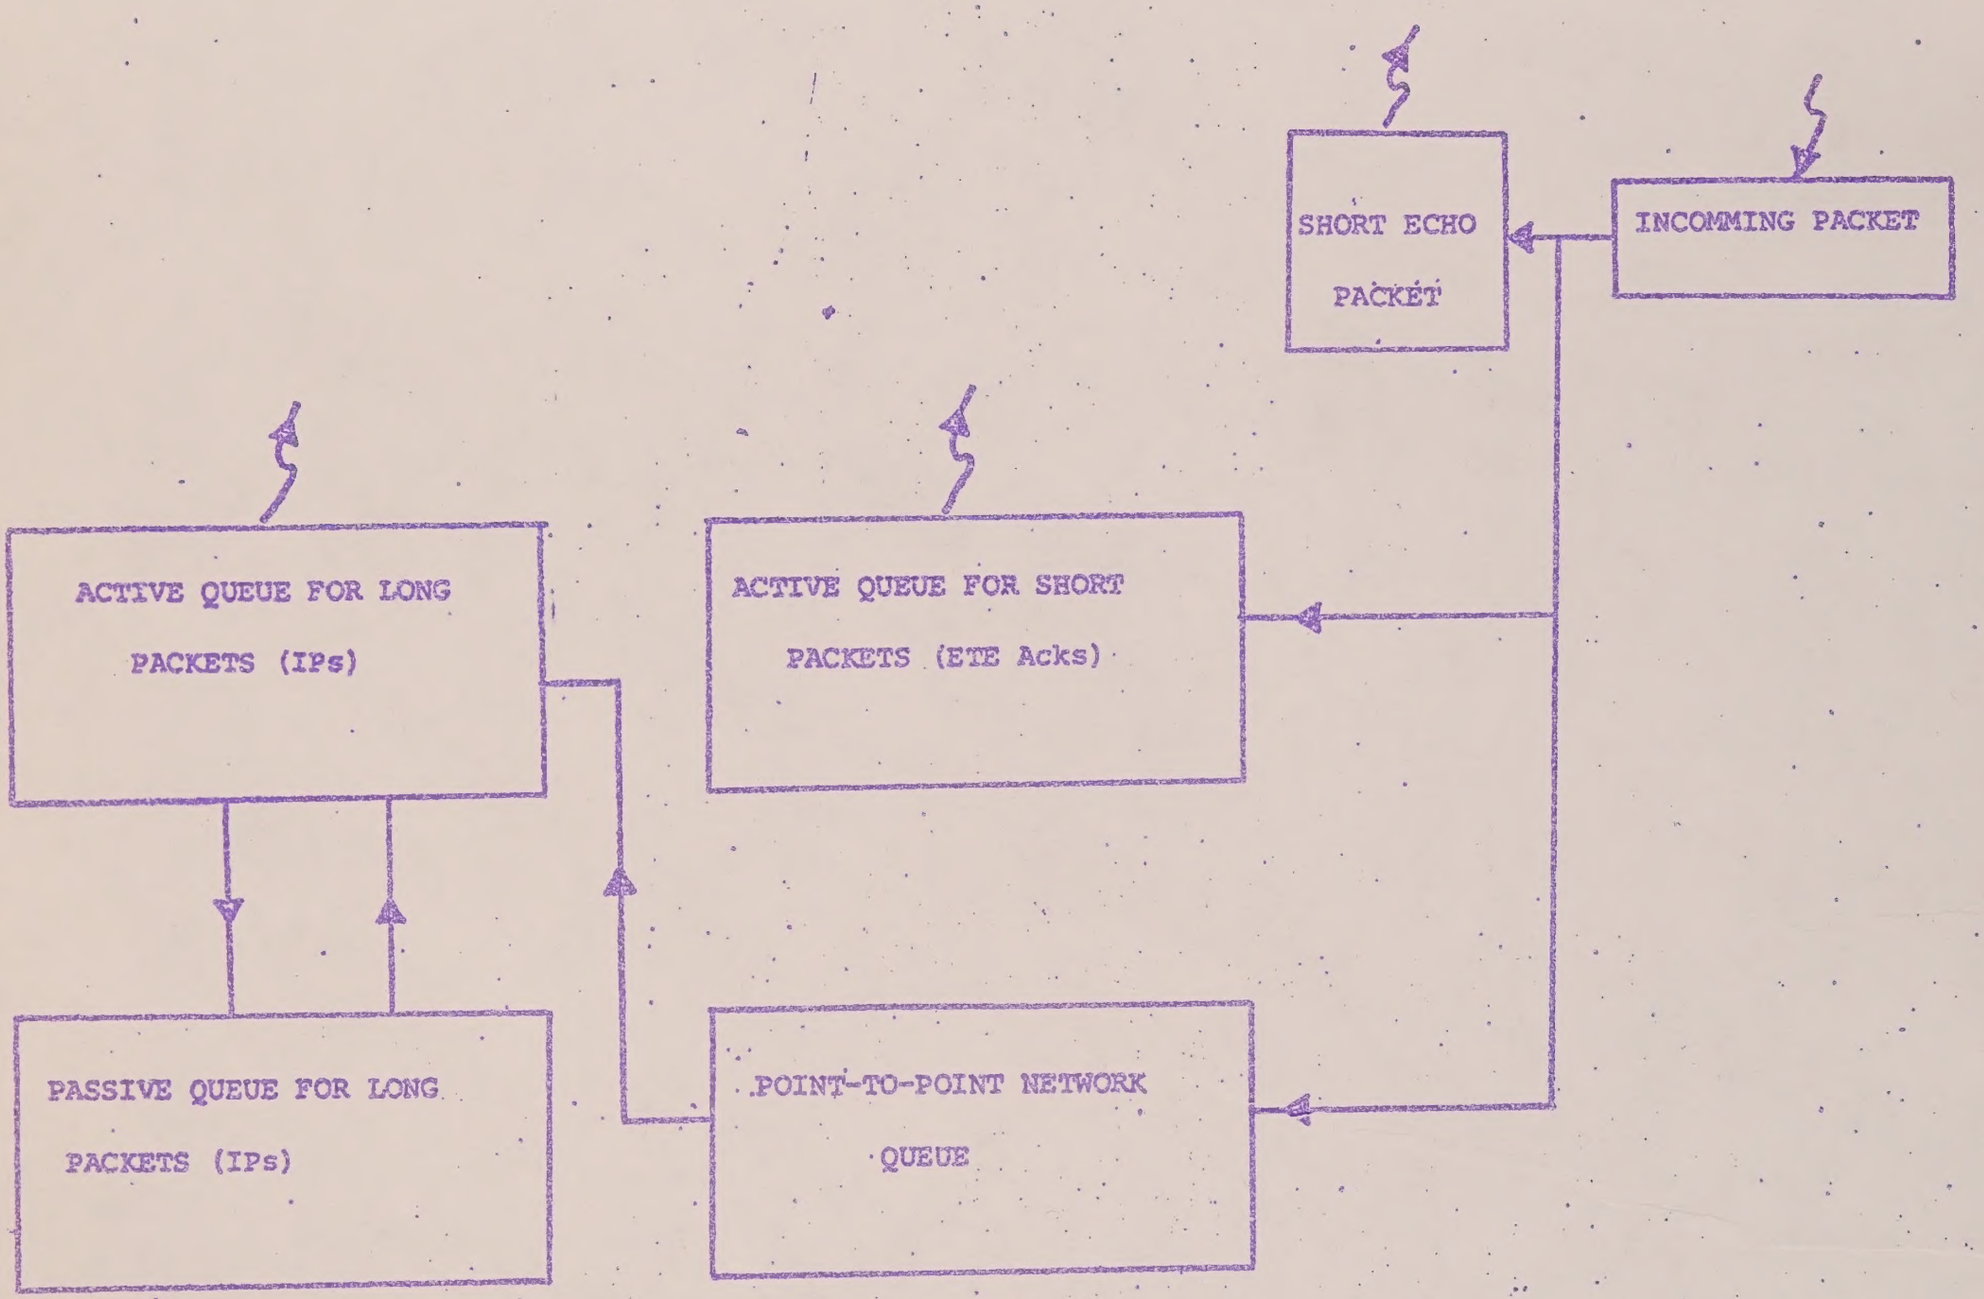
SHORT ECHO PACKET
INCOMMING PACKET
ACTIVE QUEUE FOR LONG PACKETS (IPs)
ACTIVE QUEUE FOR SHORT PACKETS (ETE Acks)
PASSIVE QUEUE FOR LONG PACKETS (IPs)
POINT-TO-POINT NETWORK QUEUE
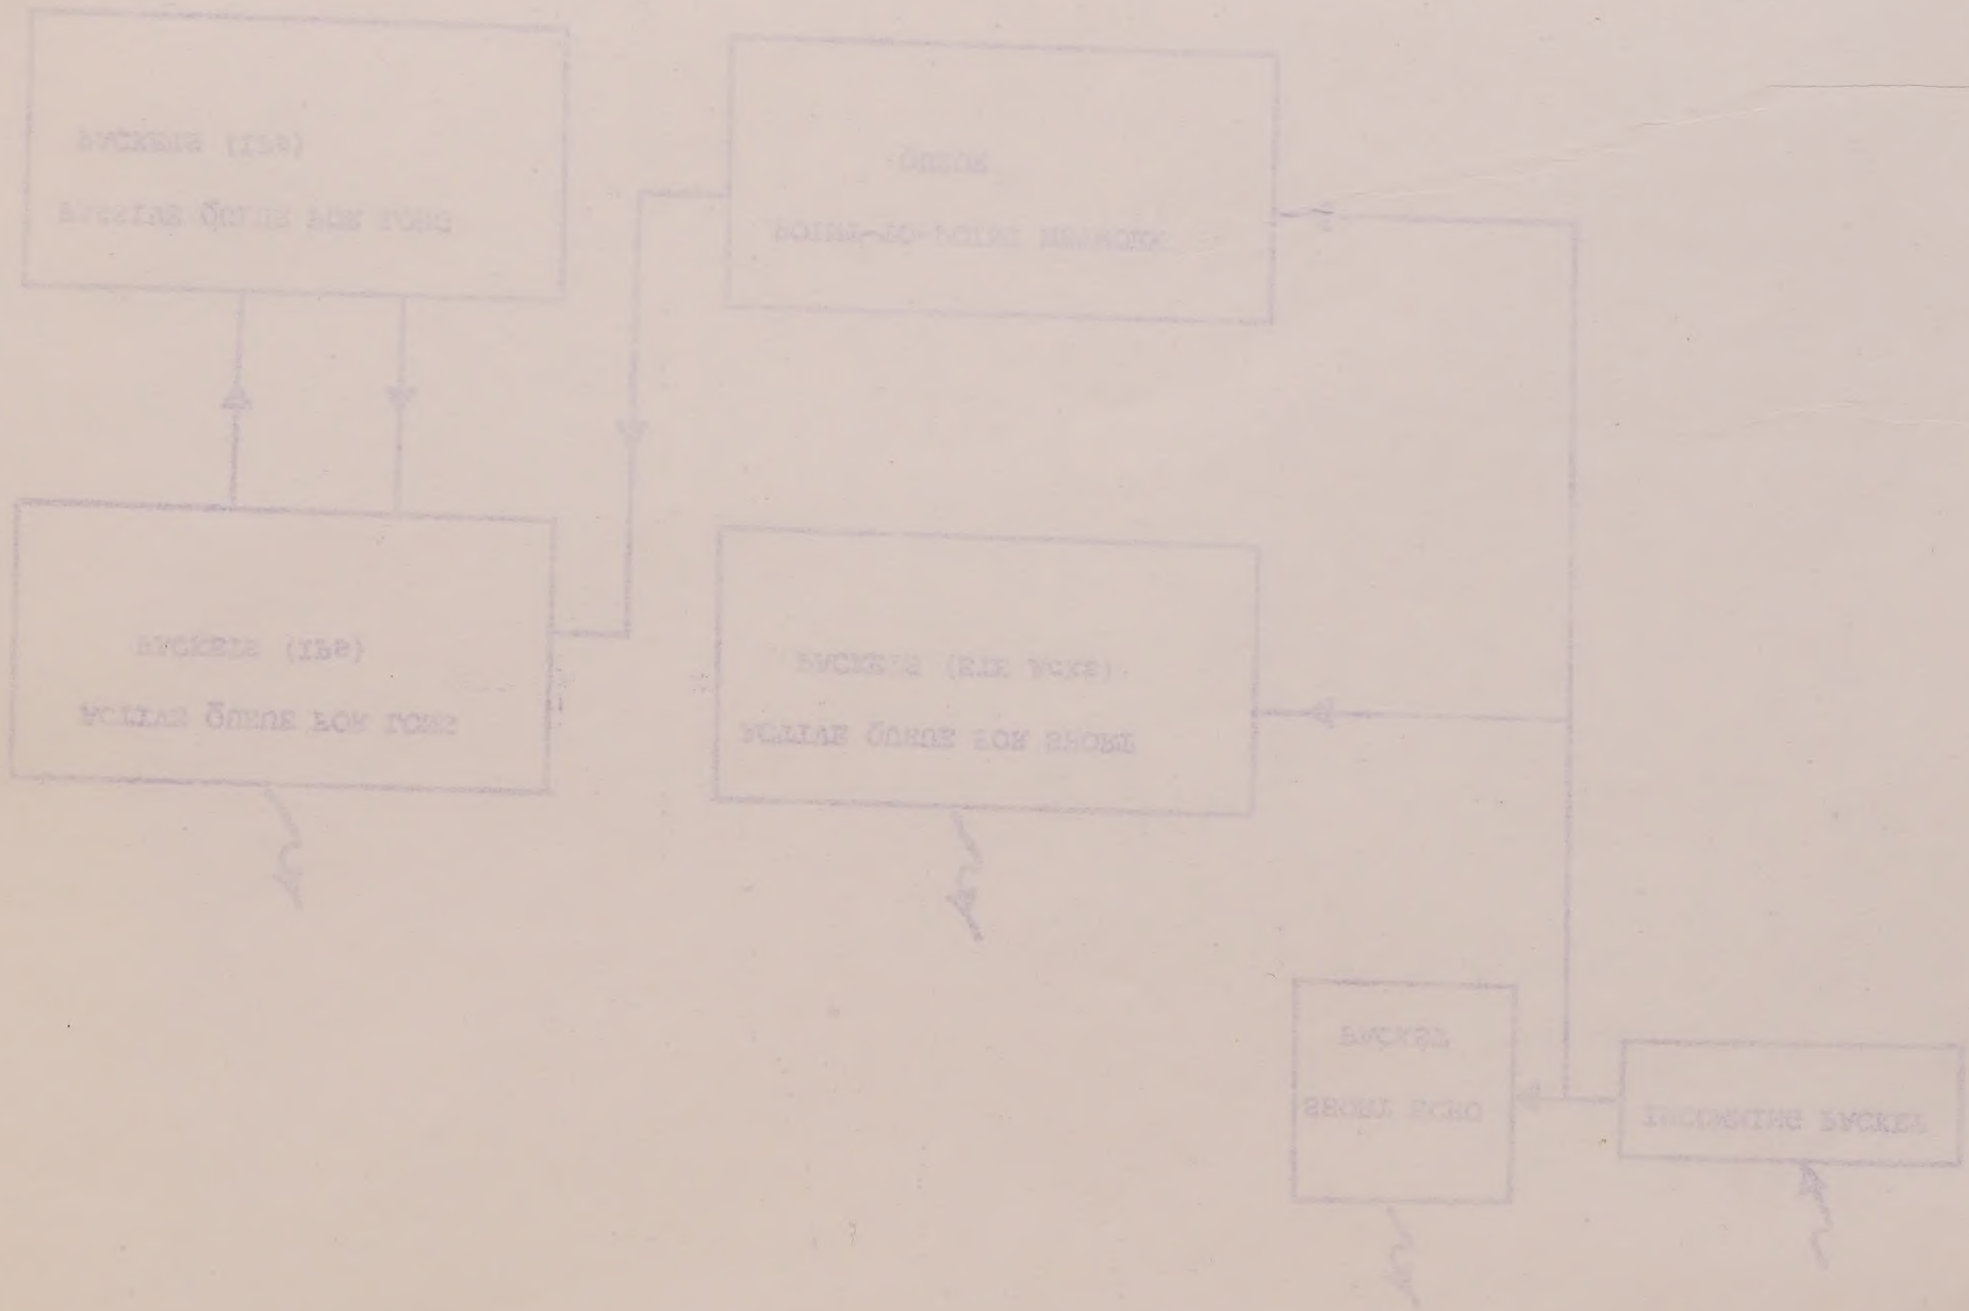
FIGURE A3: STORAGE IN STATION OF IP's AND ETE Acks.